México DF, 19 August 2016

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Requisitos de cómputo de alto rendimiento para la investigación en clima y calidad del aire

Francisco J. Doblas-Reyes
BSC-ES Director
ICREA Research Professor

Barcelona Supercomputing Center

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

- Created in 2005, now 470 employees
- Research, develop, manage information technology
- Facilitate scientific progress and its application in society

The BSC

Five departments:

- Computer Sciences
- Computer Applications for Science and Engineering
- Life Sciences
- Operations
- Earth Sciences

Hosts Marenostrum3, a 1.1 PF machine with 50,000 cores

BSC Earth Sciences Department

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

What

Environmental modelling and forecasting

Why

Our strength …
- … research …
- … operations …
- … services …
- … high resolution …

How

Develop a capability to model air quality processes from urban to global and the impacts on weather, health and ecosystems

Implement climate prediction system for subseasonal-to-decadal climate prediction

Develop user-oriented services that favour both technology transfer and adaptation

Use cutting-edge HPC and Big Data technologies for the efficiency and user-friendliness of Earth system models

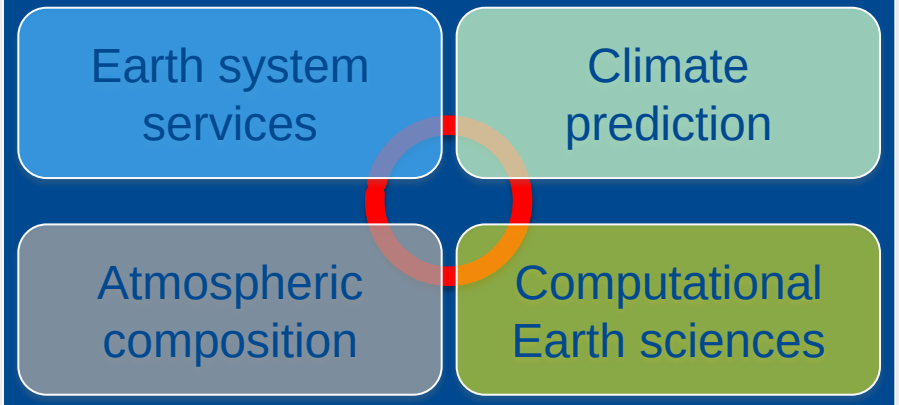

A conceptual model from CS

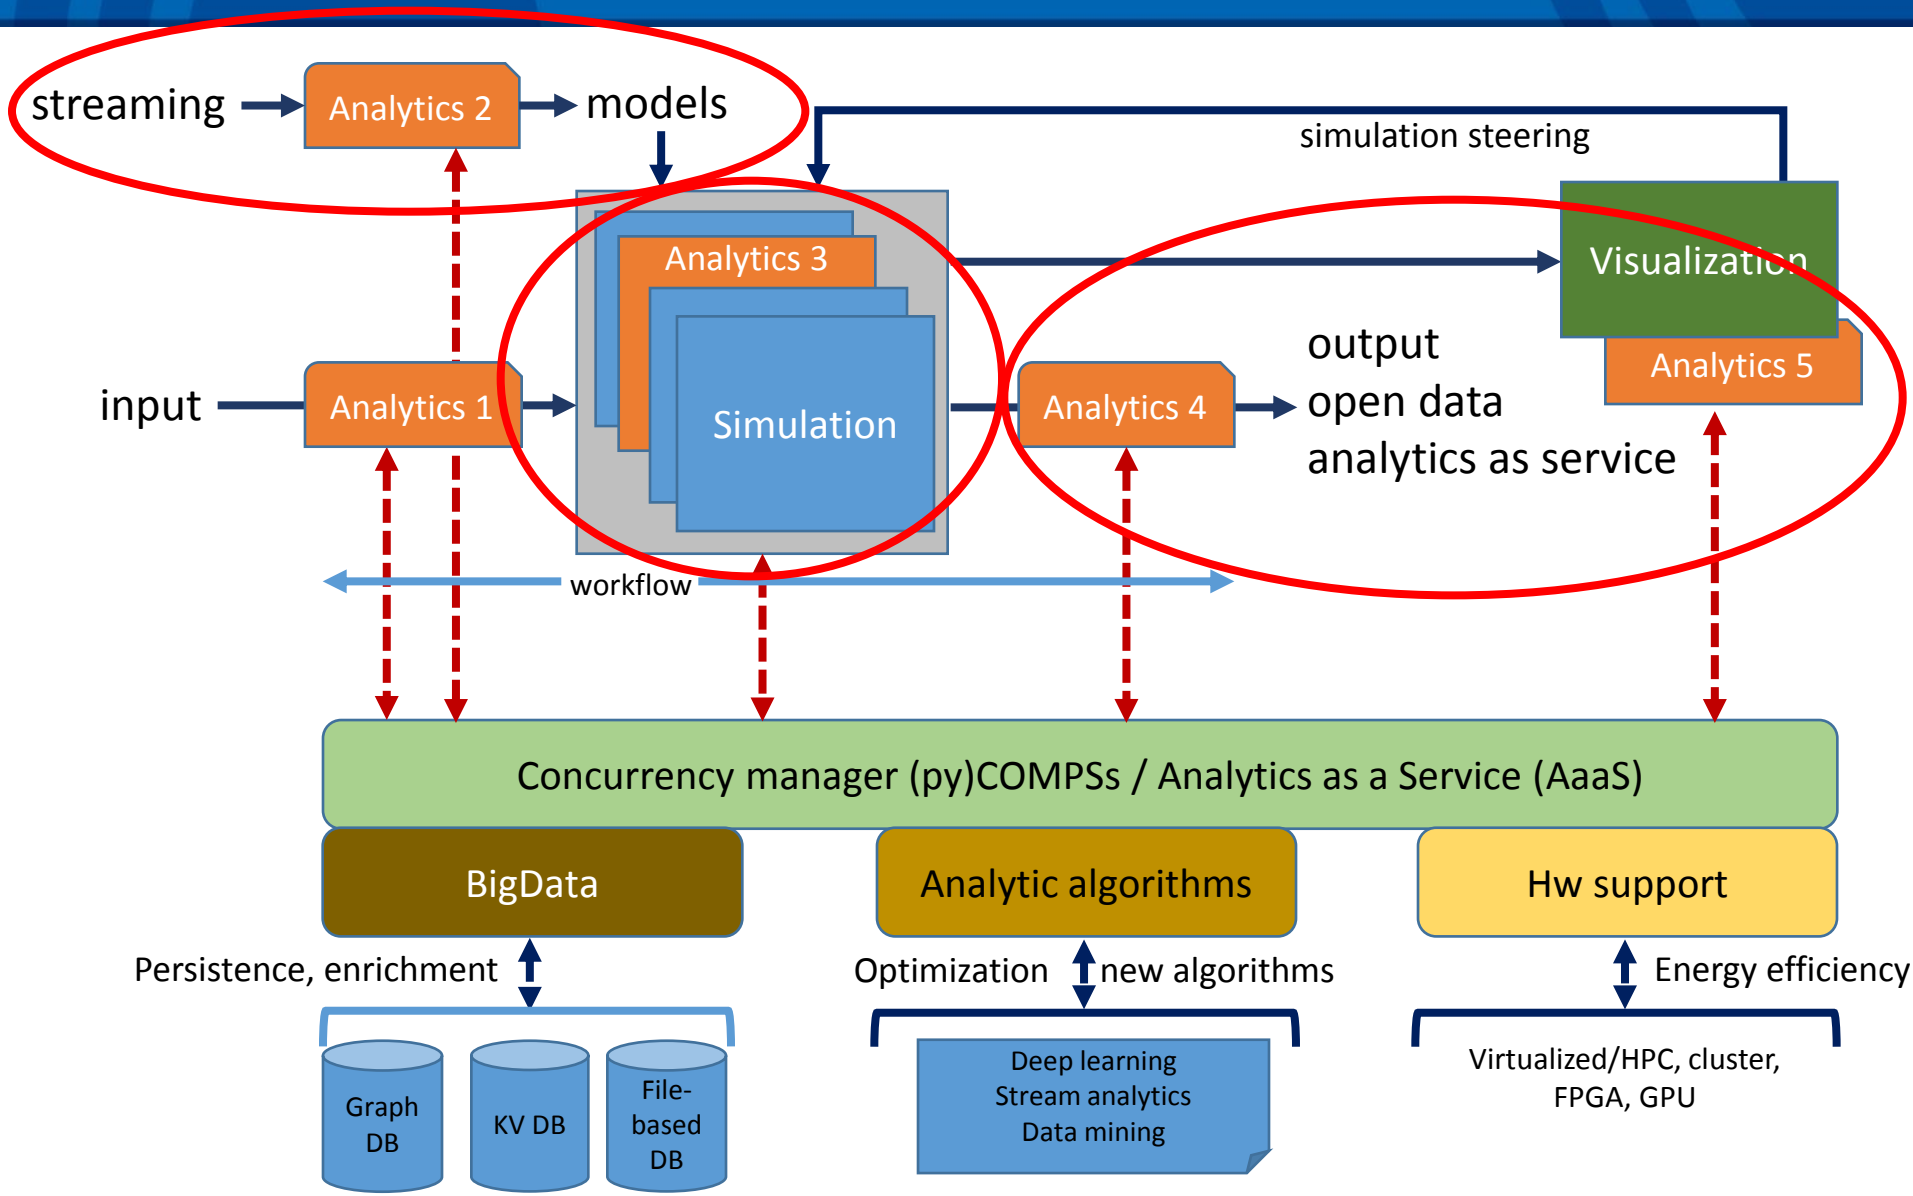

Atmospheric composition

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Atmospheric Composition

1. Development of modeling systems for atmospheric composition from local to global scales.
2. Advanced research on atmospheric composition processes using modeling techniques and Earth observations.

Air quality operational forecasts

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

AQF CALIOPE system: daily forecast and evaluation

Daily forecast for **meteorology, emissions and air quality**: Europe (12km), Iberian Peninsula (4km), Andalusia, Catalonia and Madrid (1km), since 2007

Forecast products

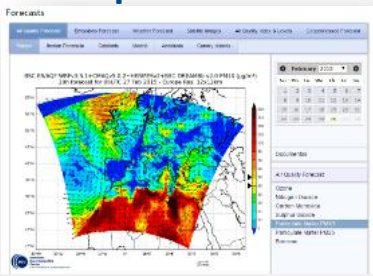
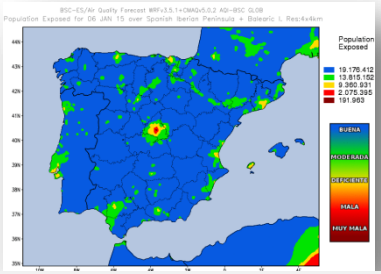

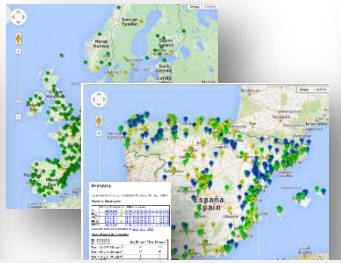
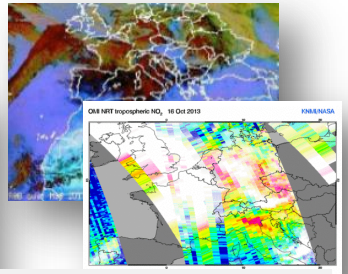

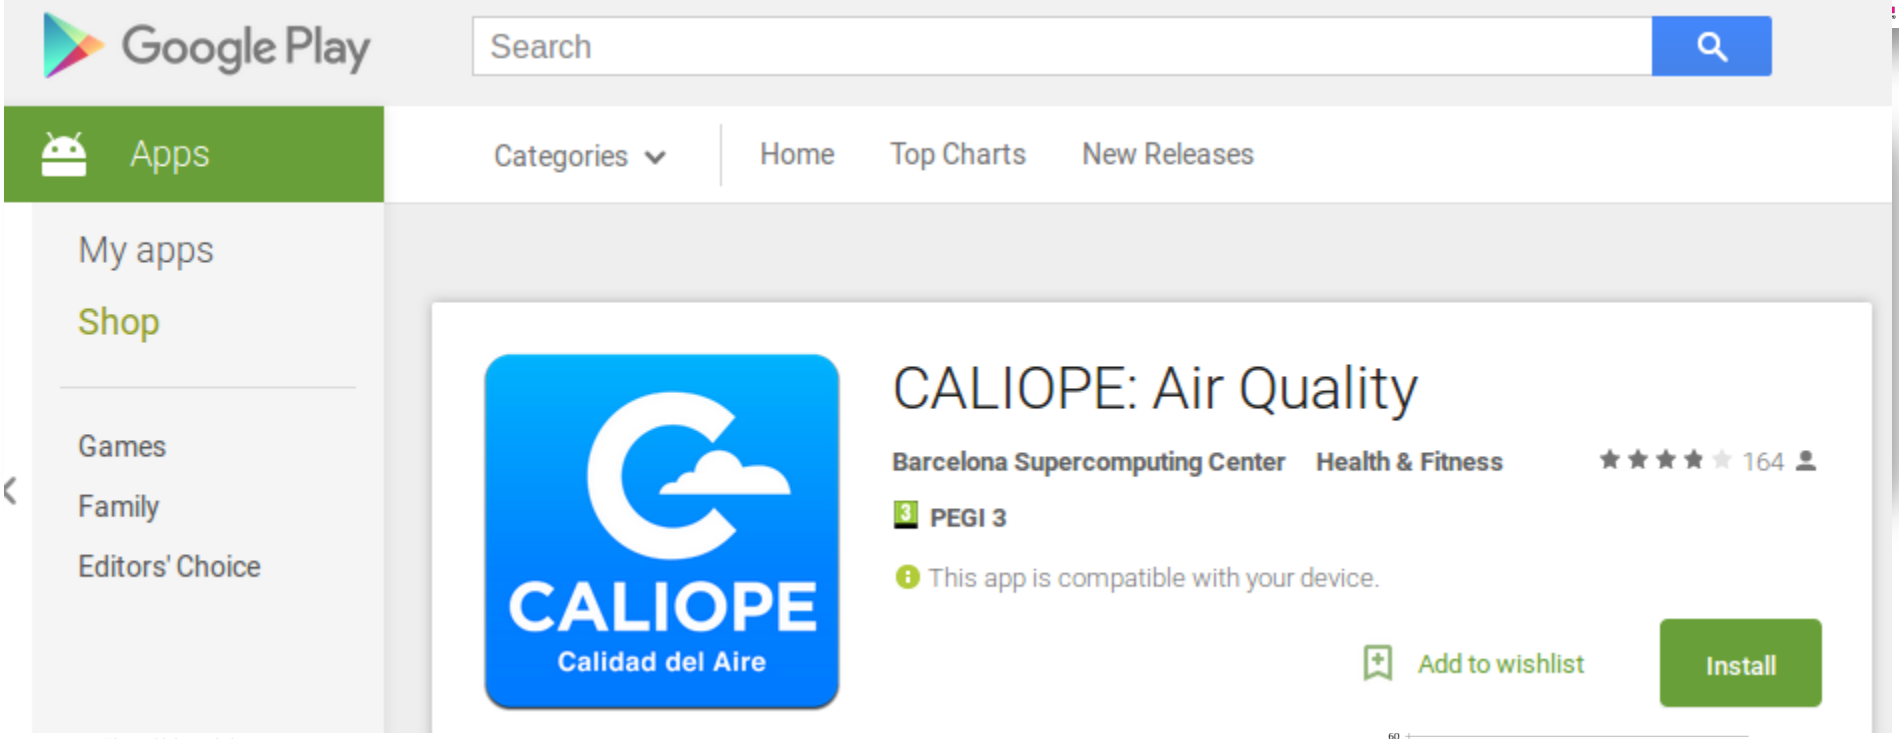

PM10_KF annual average skill evolution (2011-2014)

Impact of our air quality forecasts

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

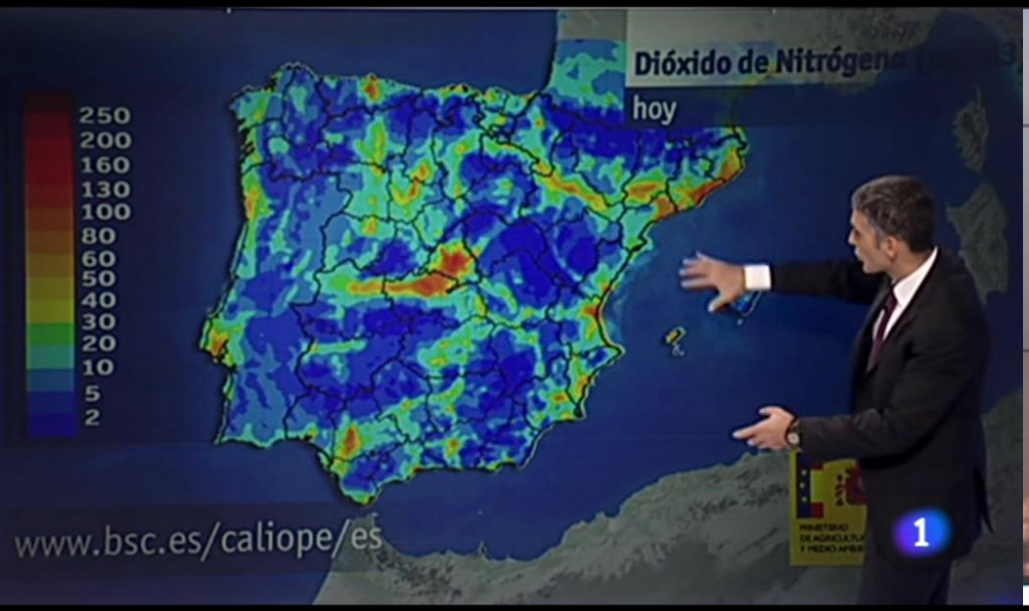

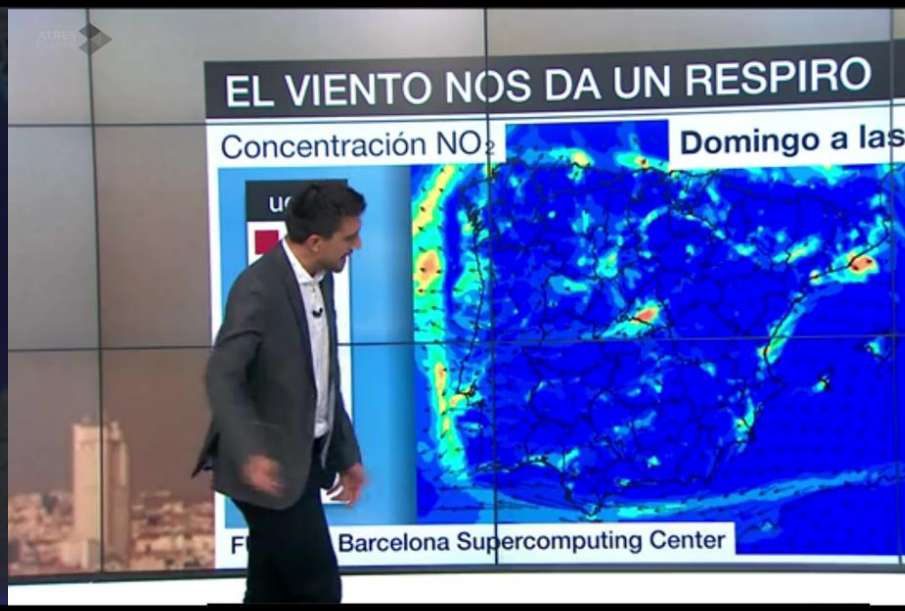

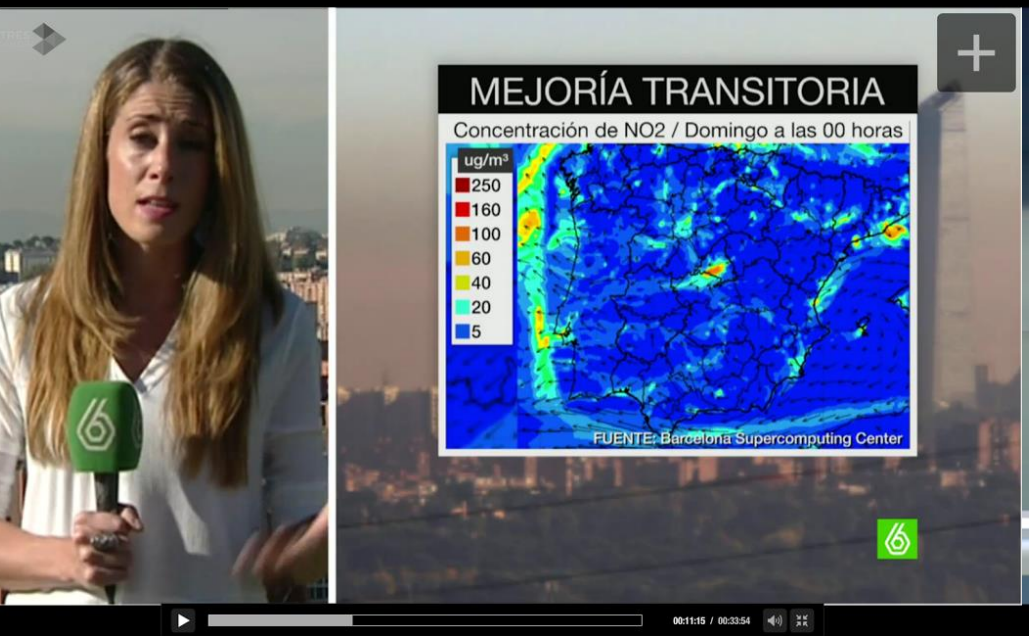

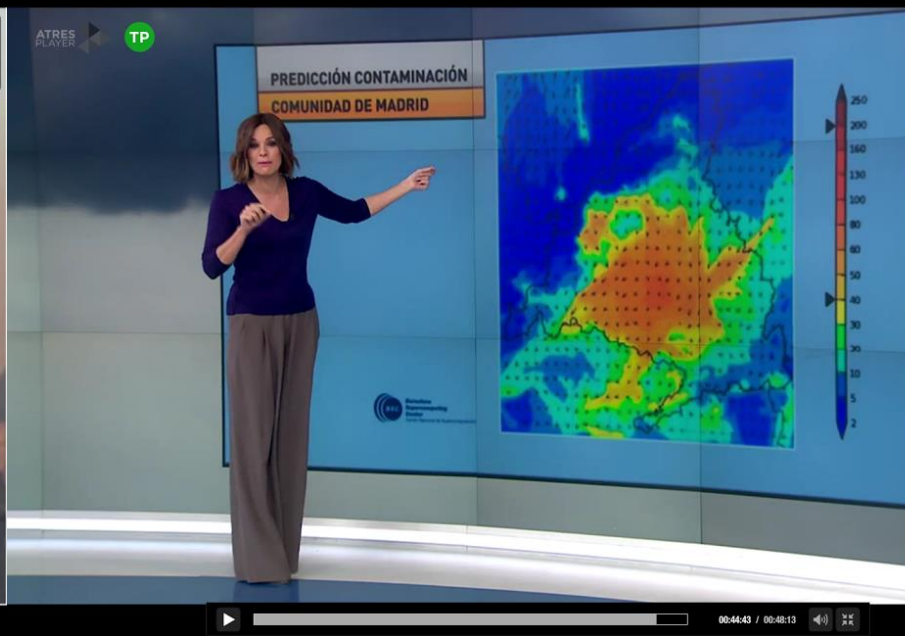

Air pollution in Mexico DF

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Meteorology

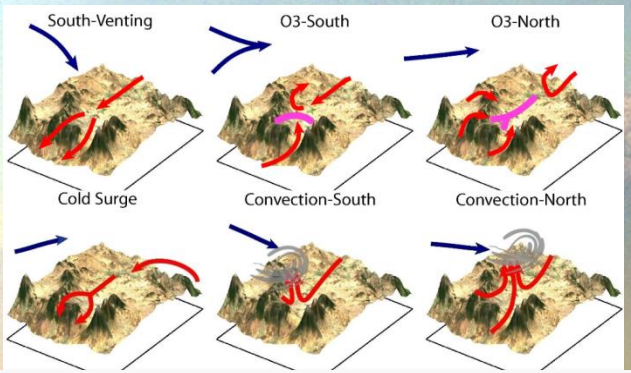

Local meteorology strongly influenced by the surrounding terrain that favours stagnant conditions and air pollution episodes

Air Quality

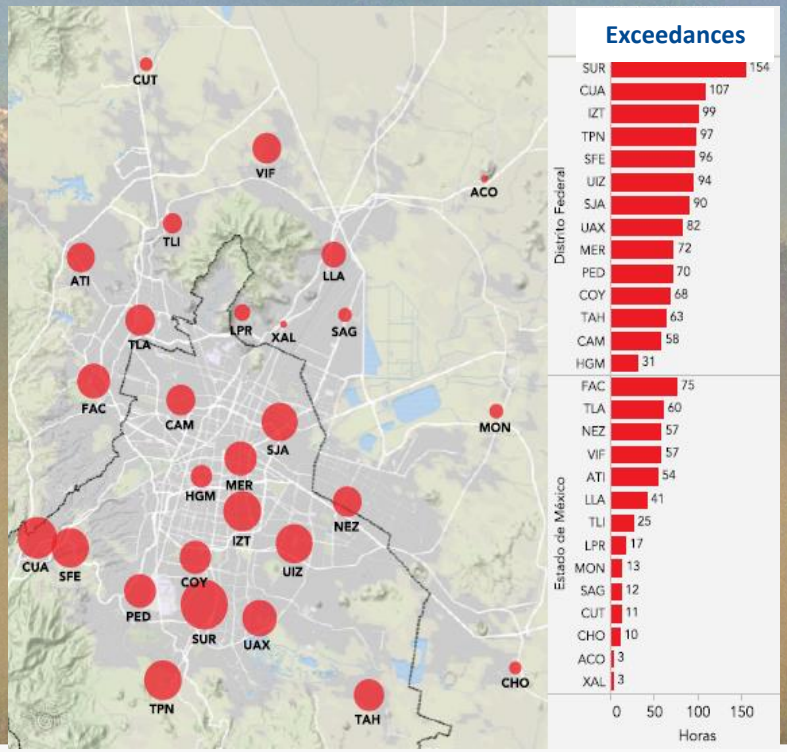

Emissions

Second largest metropolitan area in the world: more than 86,000 million vehicle-kilometre travelled per year

Exceedance of air quality limit values: Need for a management tool to develop and evaluate emission mitigation measures

Mexico City limits private car use to battle pollution

Despite criticism, authorities take drastic measures to tackle the city's first environmental emergency in 14 years.

ALJAZEERA

Air pollution in Mexico DF

Current Status

Characterization of the air quality state through a network of automatic air quality stations

Development of an air quality forecast System for Mexico DF (AQFS-MexDF), a modelling tool for air quality management:

- Complement the public information service provided by the monitoring network
- Know in advance the possibility that air pollution episodes occur
- Contribute to the development and evaluation of air quality plans (ProAire)

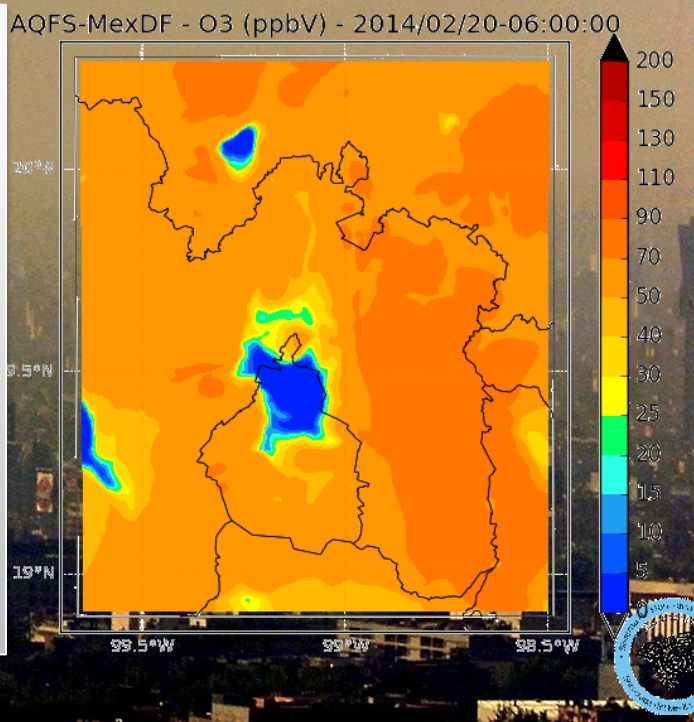

Unified modelling solution

NMMB/BSC-CTM: Unified multiscale atmosphere-chemistry modelling system across spatial scales.

Joint venture with NCEP.

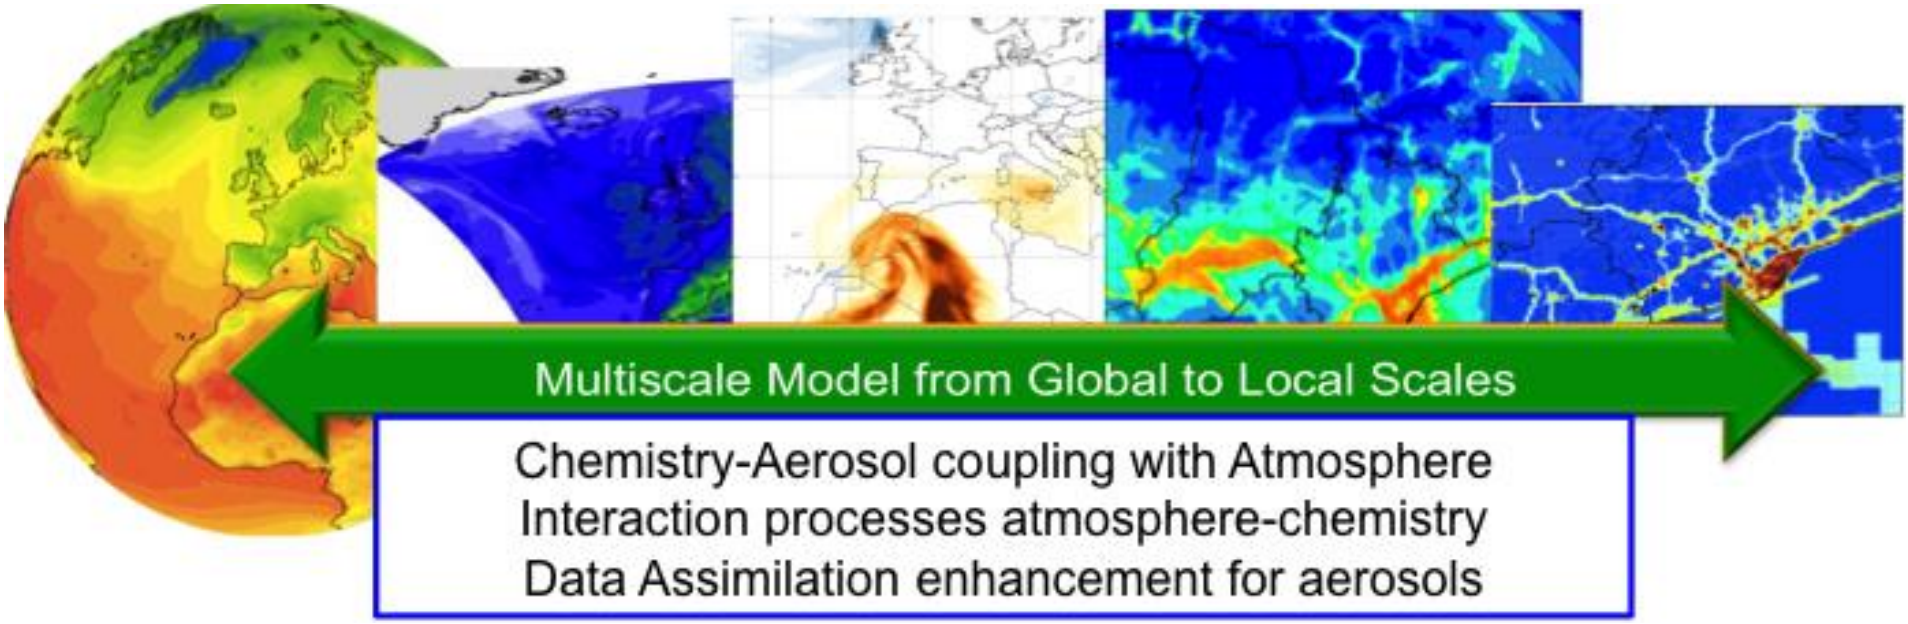

New air quality management tools

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Where we are now

Where we want to be

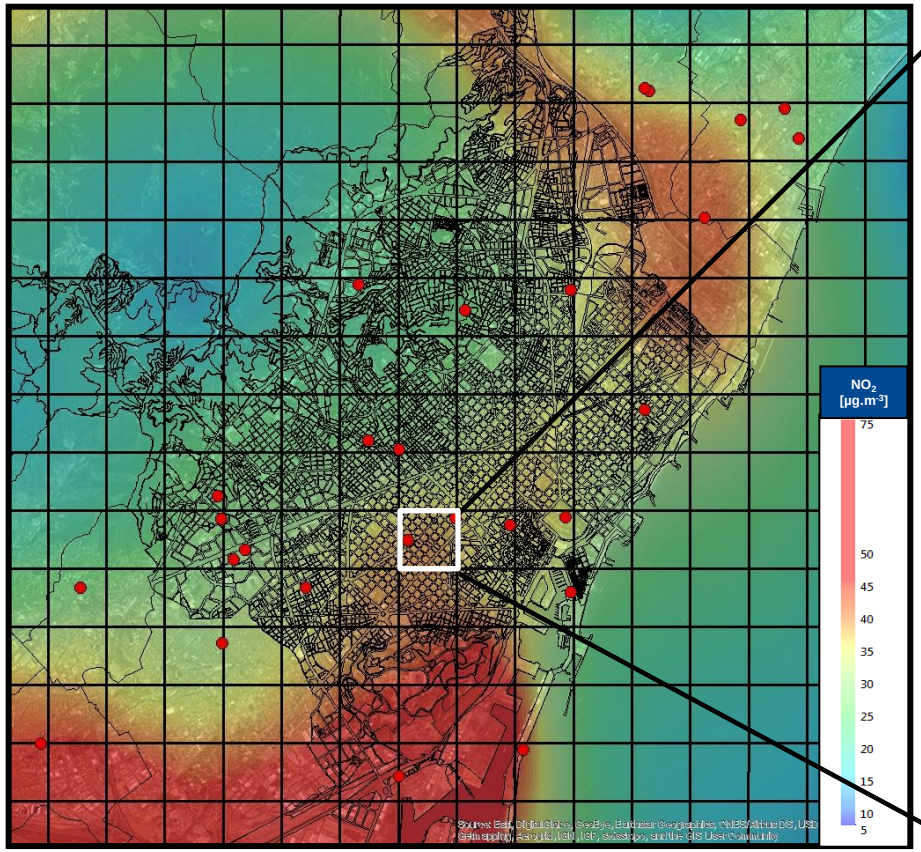

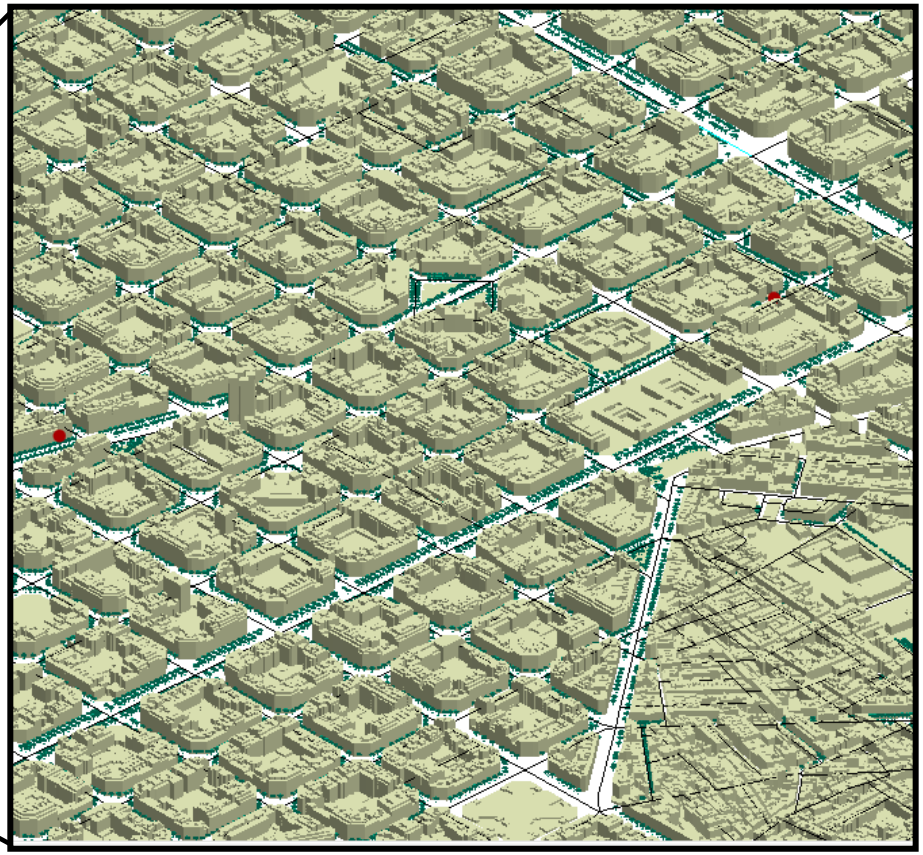

Objective	To develop an air quality model based on a CFD coupled to an atmospheric chemistry model at city scale, enhanced by the use of Big Data technologies, to assess urban air quality.

Data streaming for emissions

High-resolution air quality modelling requires appropriate emissions

- Collection and processing of sensor-generated data to feed a real-time emission model (and to validate air quality predictions)
- Providing the sensors with the adequate technology is a challenge
- Managing large volumes is another one: sampling 10 Hz, ten variables ~30 MB/day, city-wide ~300 GB/day (10,000 vehicles)

Many more sensors possible

Need for computing before transmission

https://envirocar.org/

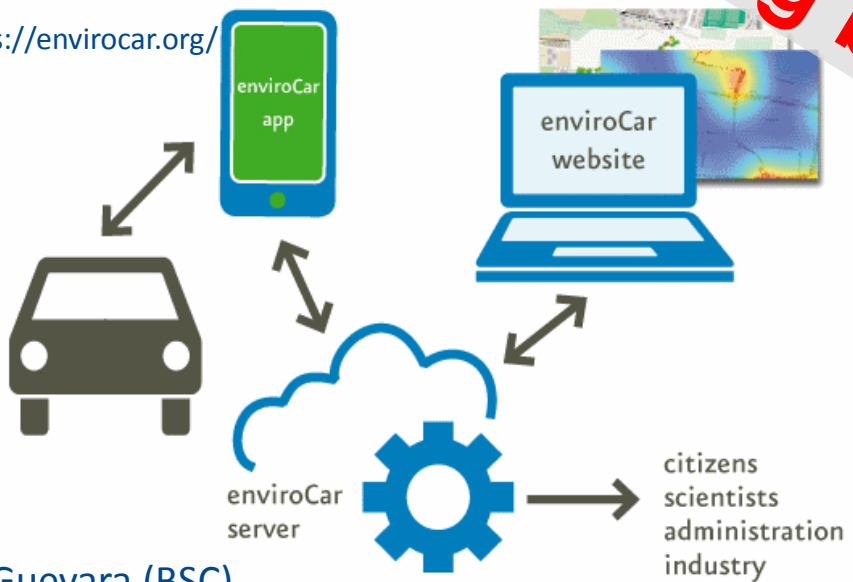

CO concentrations

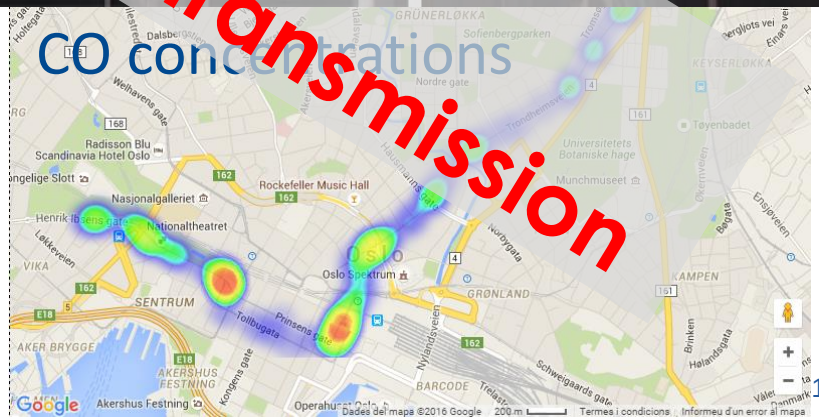

M. Guevara (BSC)

The WMO dust forecast centres

NMMB/BSC-CTM is used for, among many other things, producing operational dust forecasts.

MODEL	RUN TIME	DOMAIN	DATA ASSIMILATION
BSC-DREAM8b	12	Regional	No
MACC	00	Global	MODIS AOD
DREAM-NMME-MACC	12	Regional	MACC analysis
NMMB/BSC-Dust	12	Regional	No
MetUM	00	Global	MODIS AOD
GEOS-5	00	Global	MODIS reflectances
NGAC	00	Global	No
EMA REG CM4	12	Regional	No
DREAMABOL	12	Regional	No

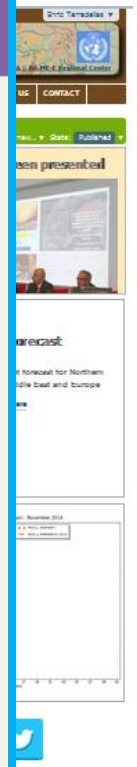

WMO Sand
Advisory and Assessment System
Regional Center for Northern Africa, Middle East and Europe http://sds-was.aemet.es

ter
http://dust.aemet.es

The WMO dust forecast centres

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

SDS web site visits 1 Jan 2015-12 Mar 2015

SDS-WAS NAMEE Regional Center

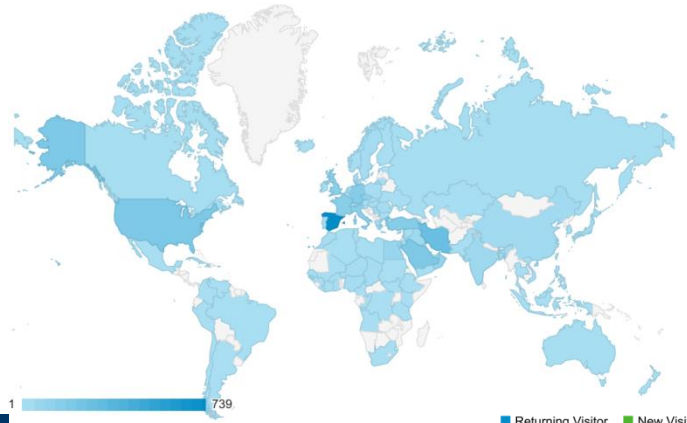

Barcelona Dust Forecast Center

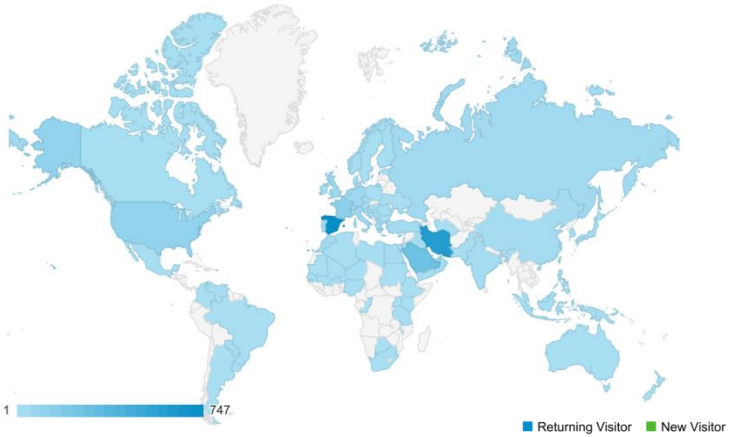

GOBIERNO DE ESPAÑA
MINISTERIO DE MEDIO AMBIENTE Y MEDIO RURAL Y MARINO
AEMet Agencia Estatal de Meteorología

Barcelona Dust Forecast Center

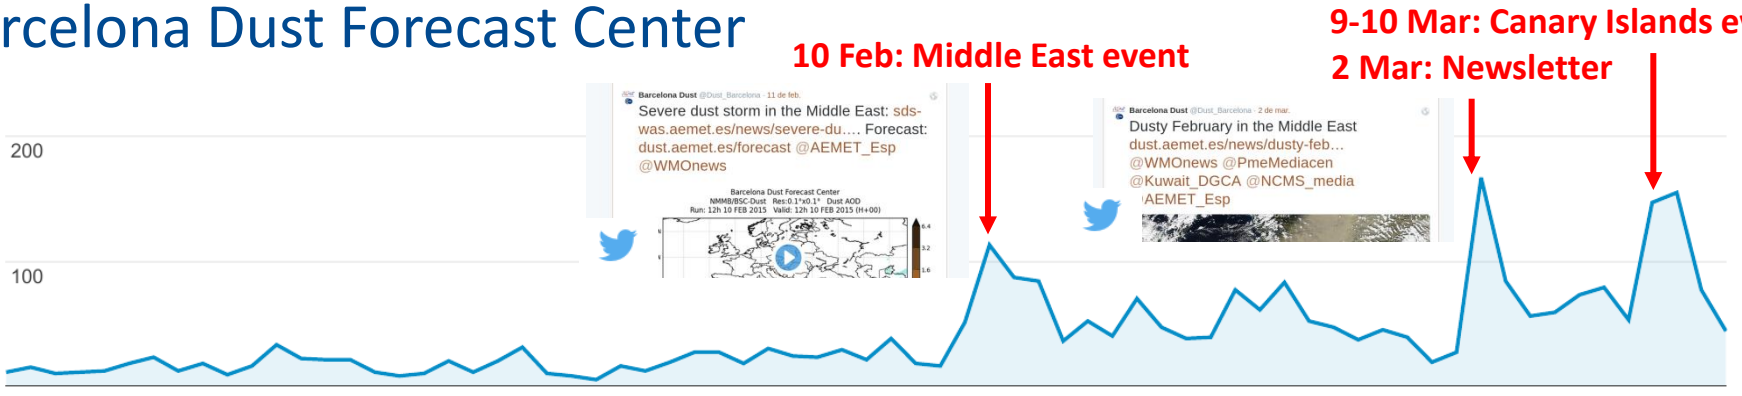

Climate prediction

Climate Prediction

1. Development of climate prediction systems focused on sub-seasonal to decadal time scales.
2. Exploitation of climate predictions to understand predictability mechanisms and forecast system limitations.

Climate prediction time scales

Progression from initial-value problems with weather forecasting at one end and multi-decadal to century projections as a forced boundary condition problem at the other, with climate prediction (sub-seasonal, seasonal and decadal) in the middle. Prediction involves initialization and systematic comparison with a simultaneous reference.

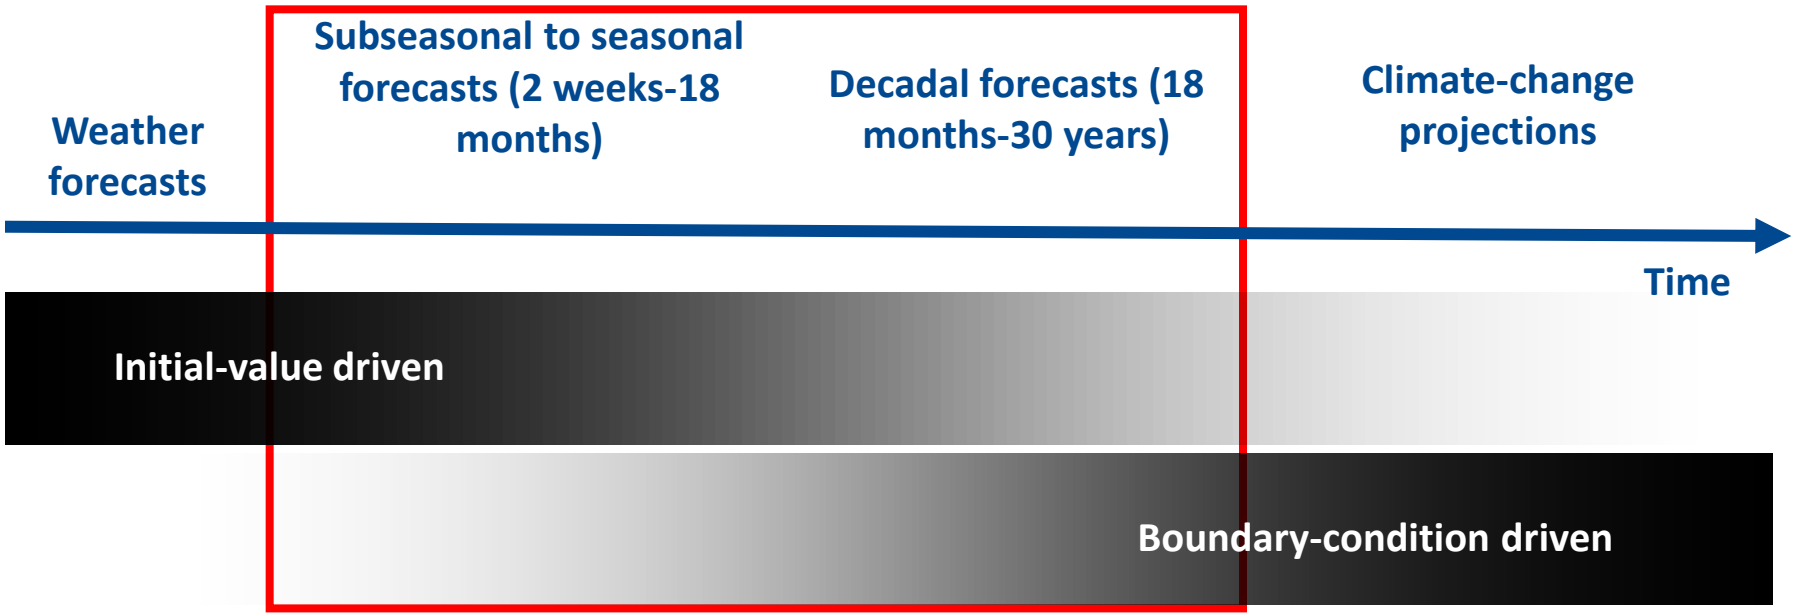

Adapted from Meehl et al. (2009)

Climate change is taking place

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Rank of the 2015 annual mean temperature over the last 37 years from ERA Interim.

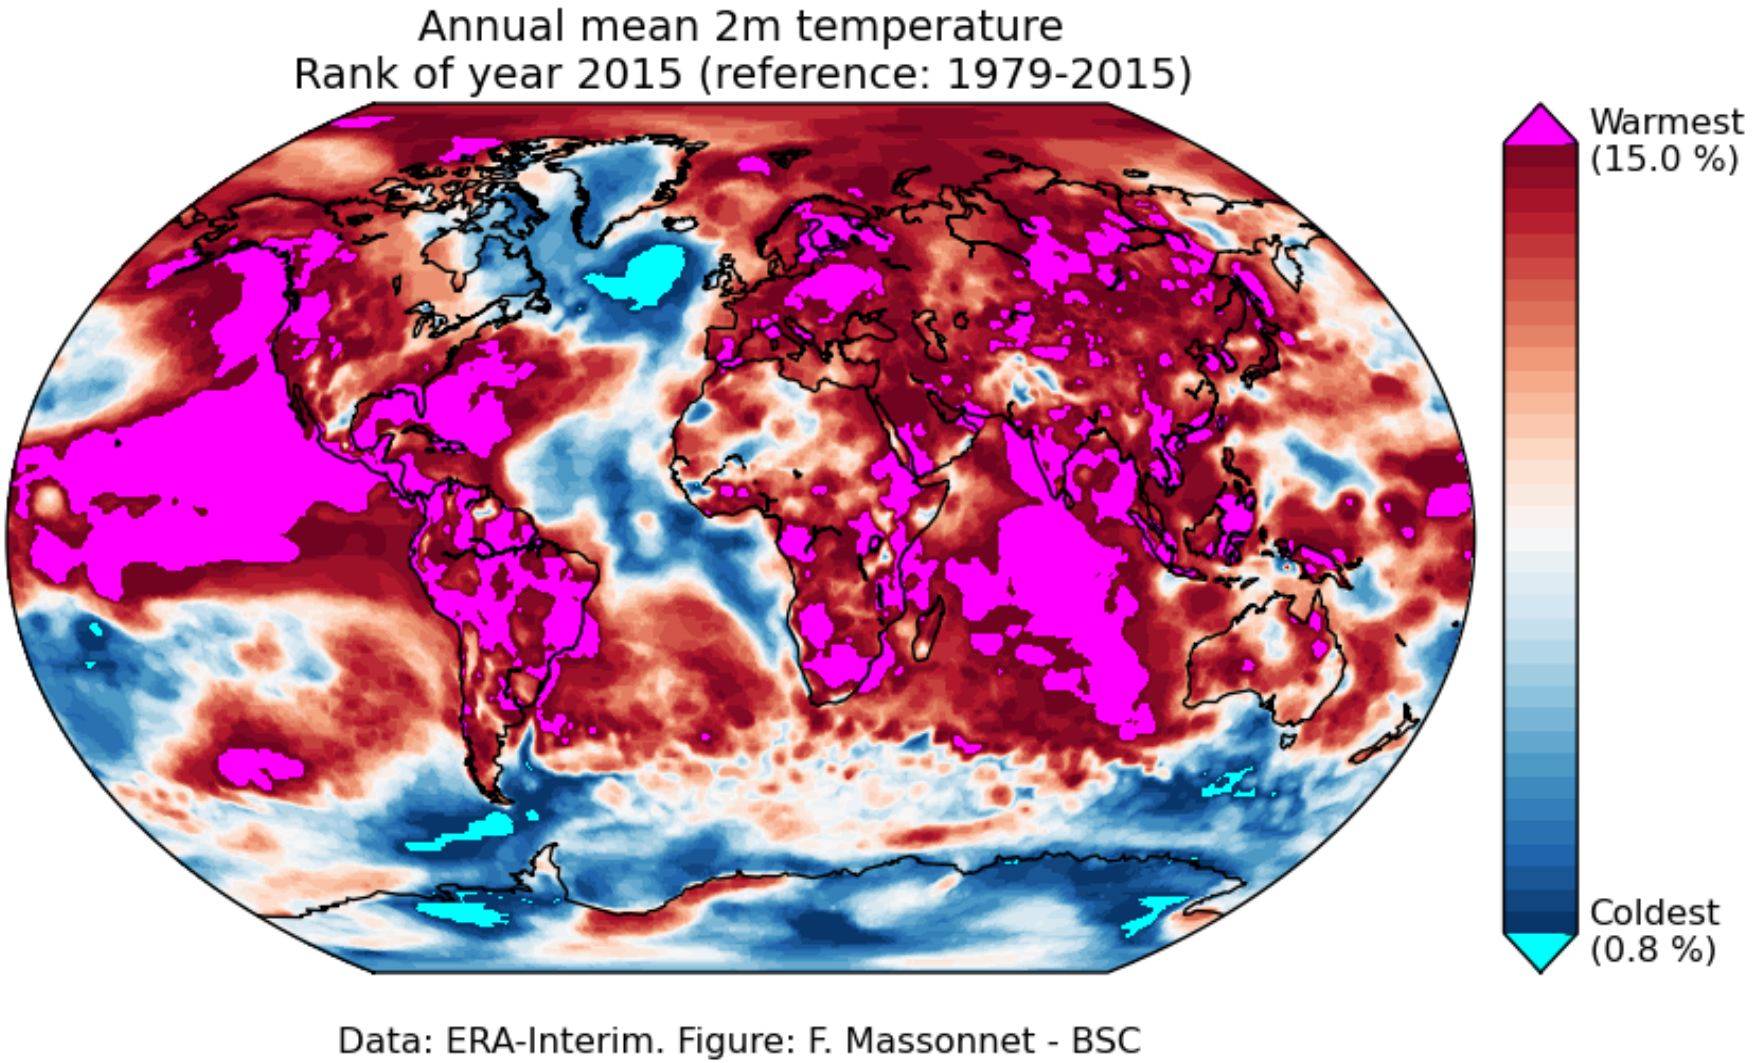

Climate predictions

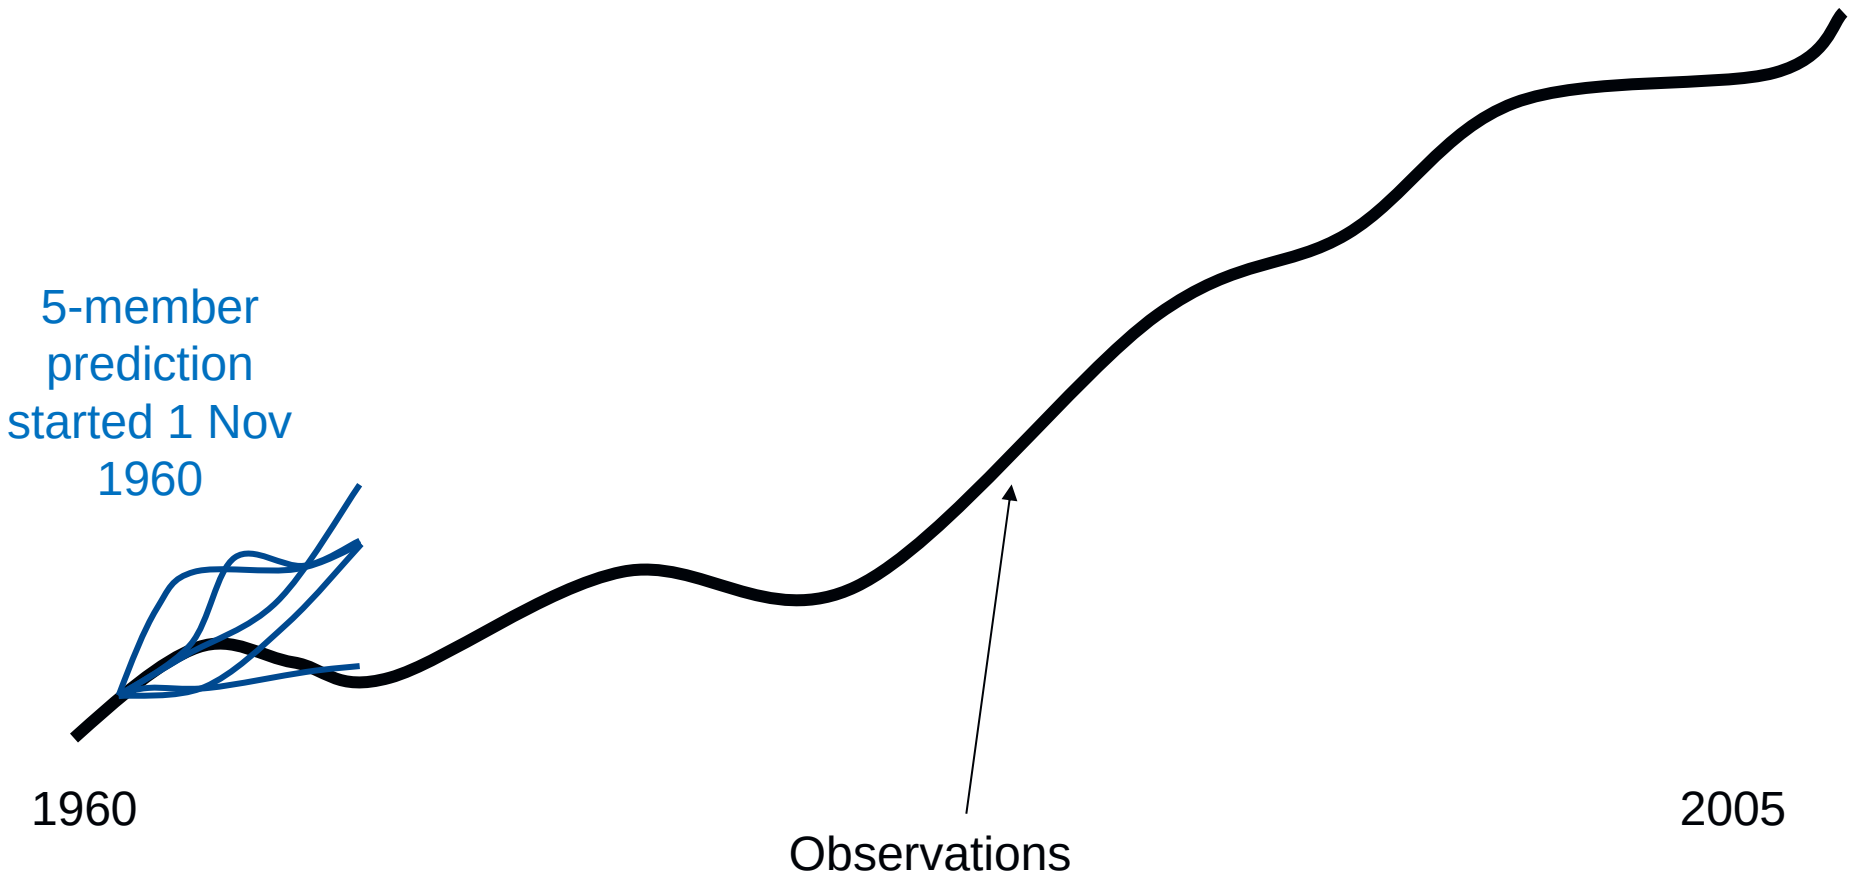

Climate predictions

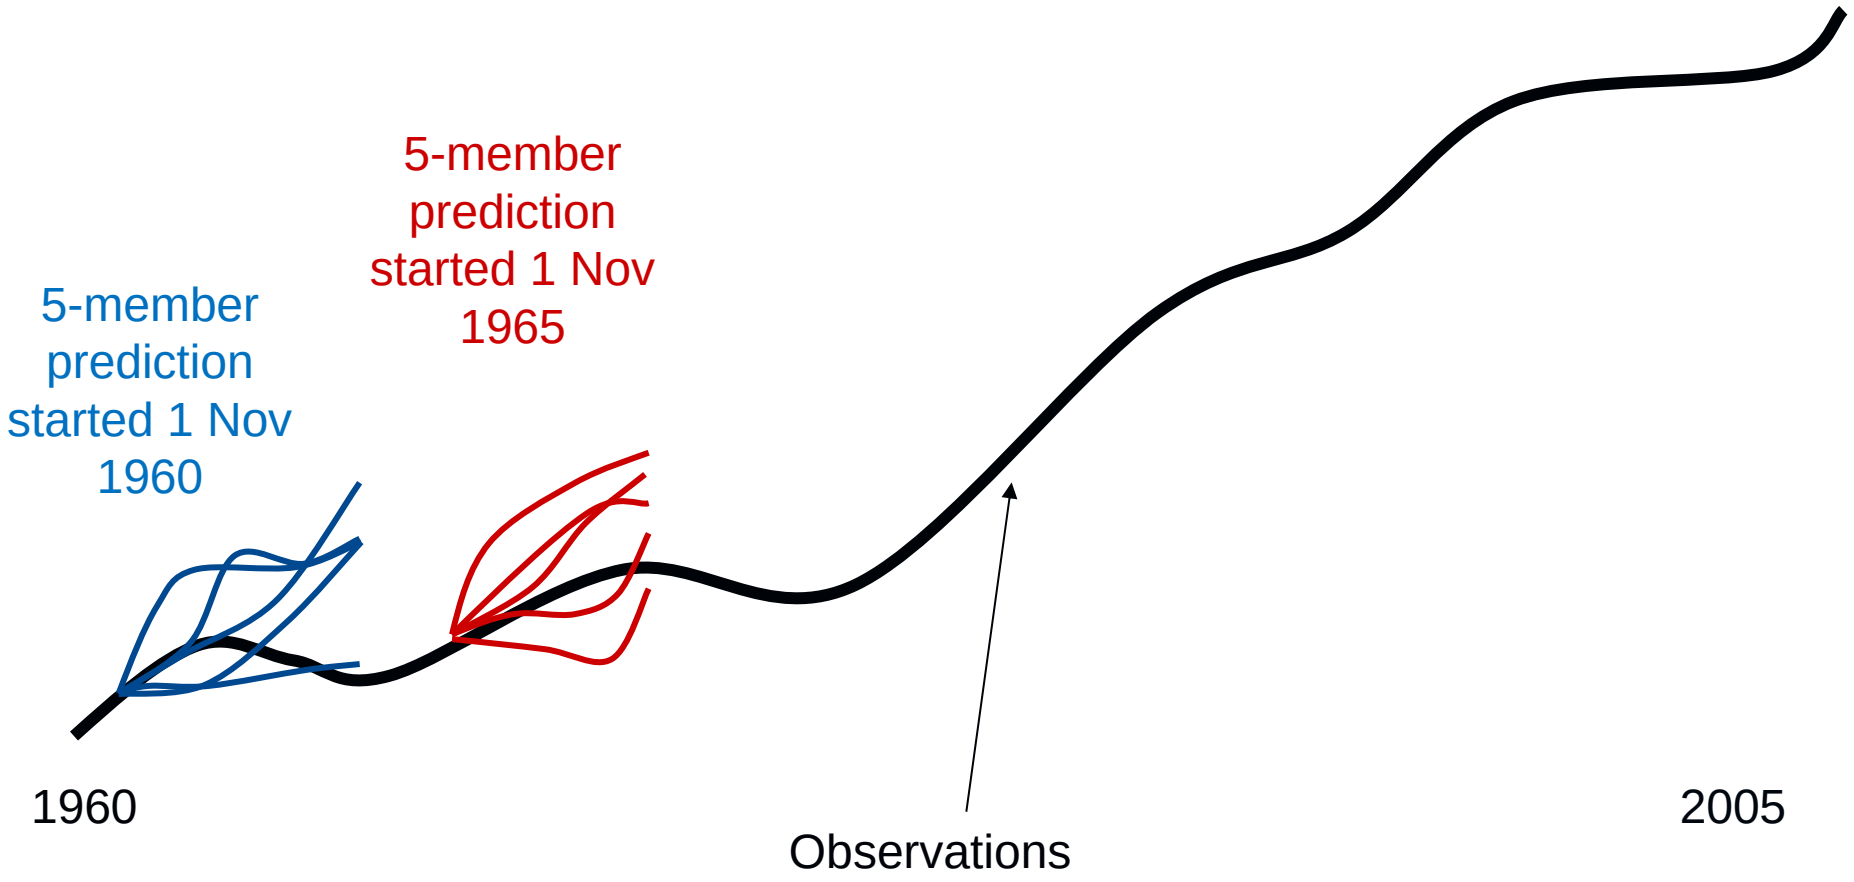

Climate predictions

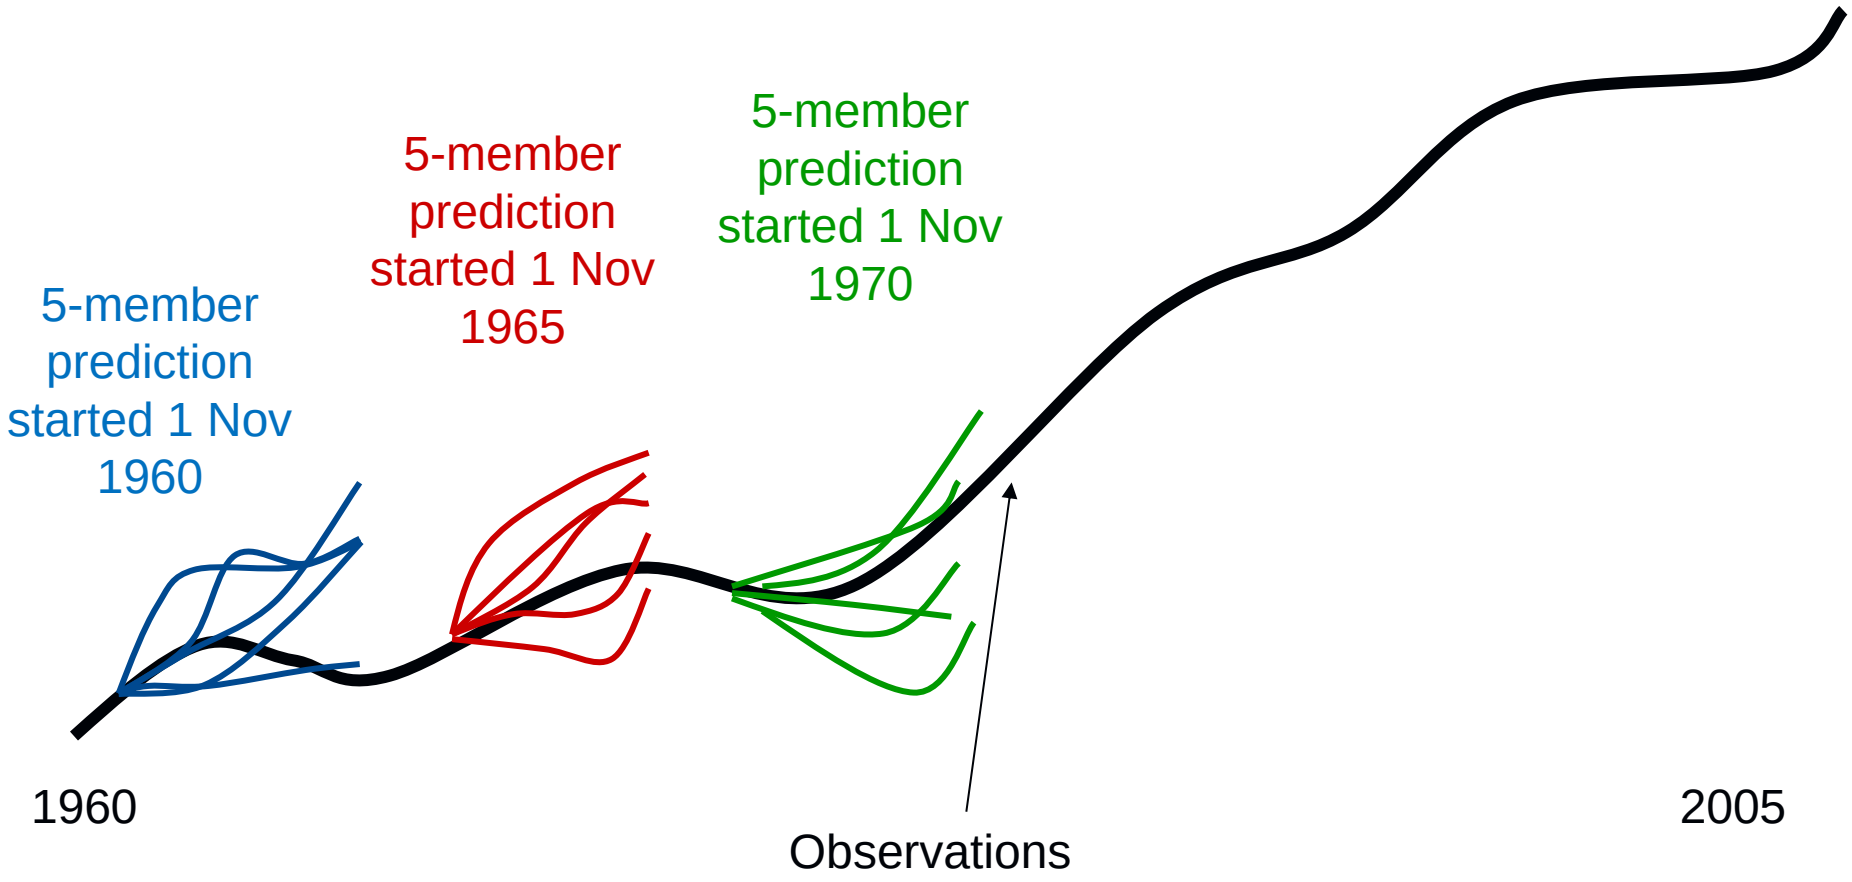

Climate predictions

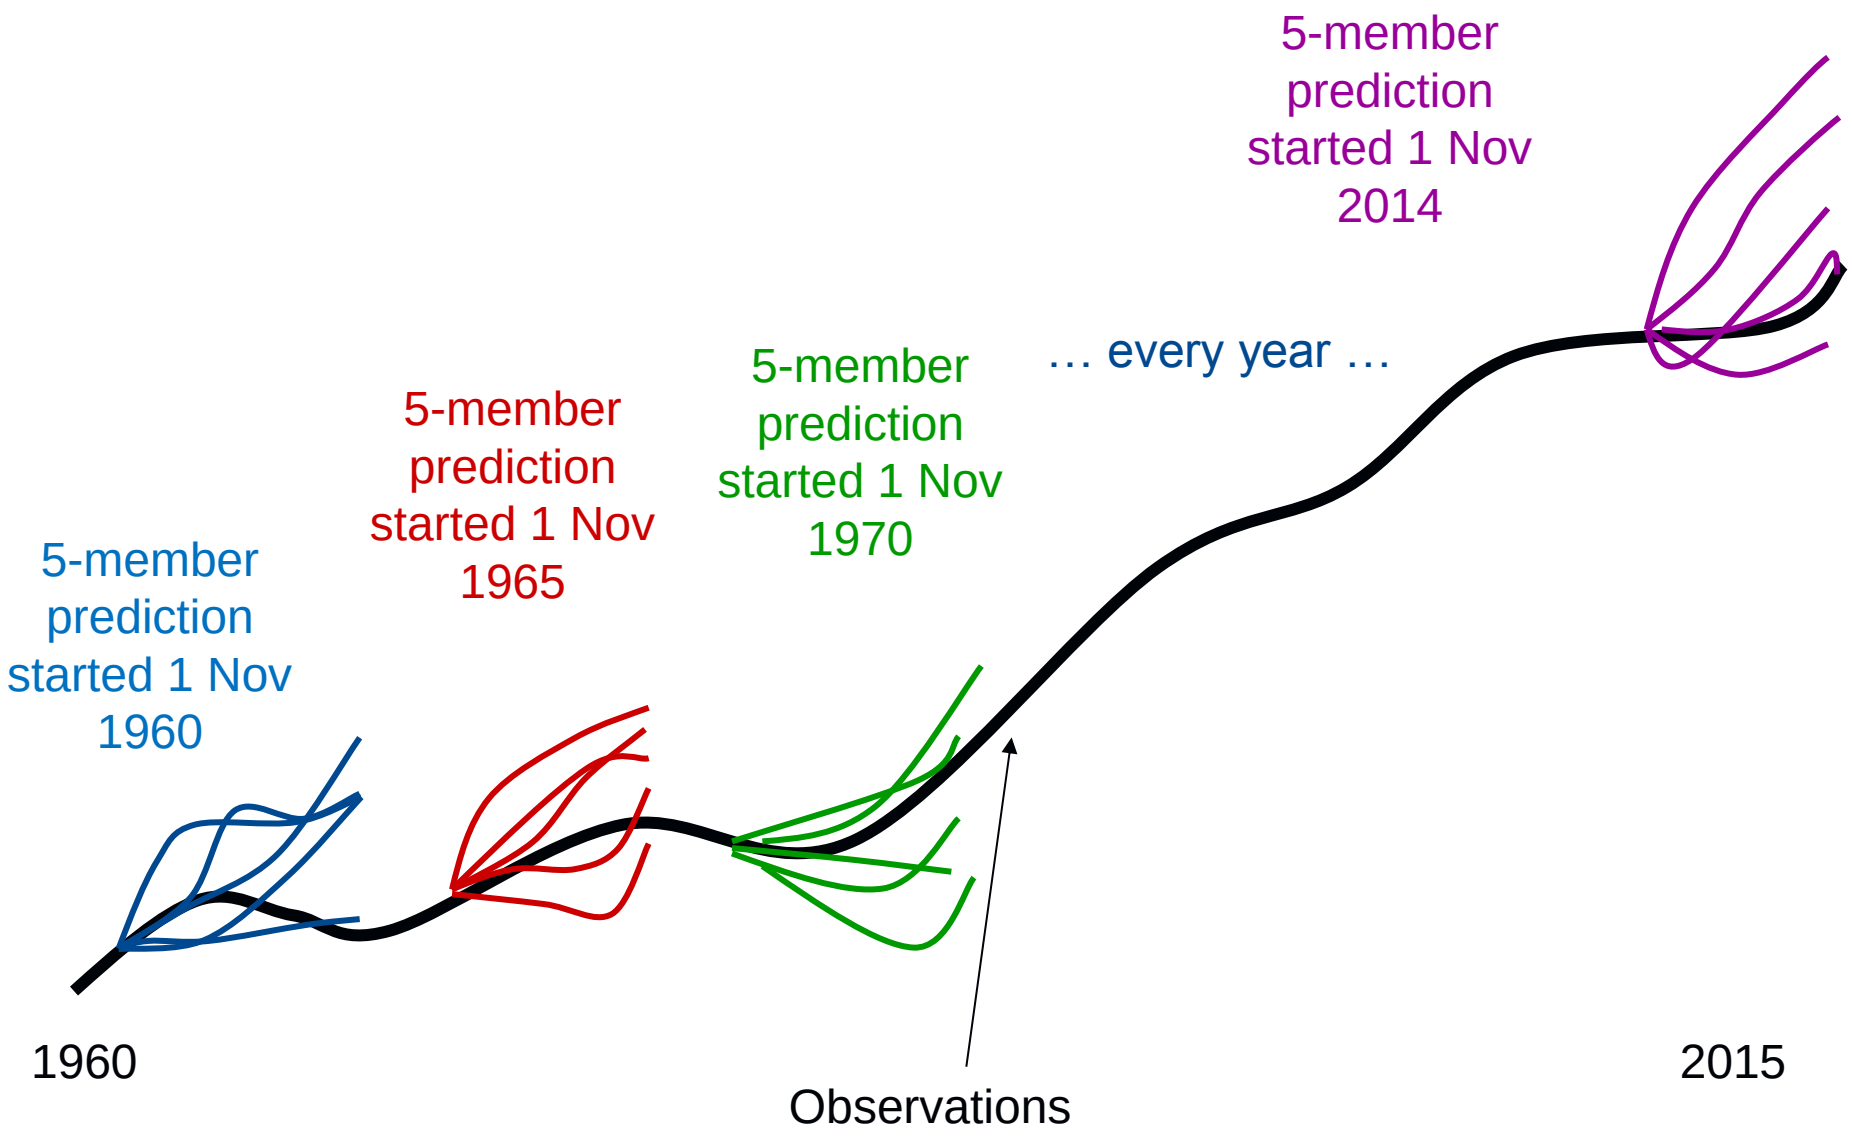

Climate predictions

Typical sizes: ten members, ten forecast years, 55 start dates -> 550 independent simulations -> 5,500 years of simulation

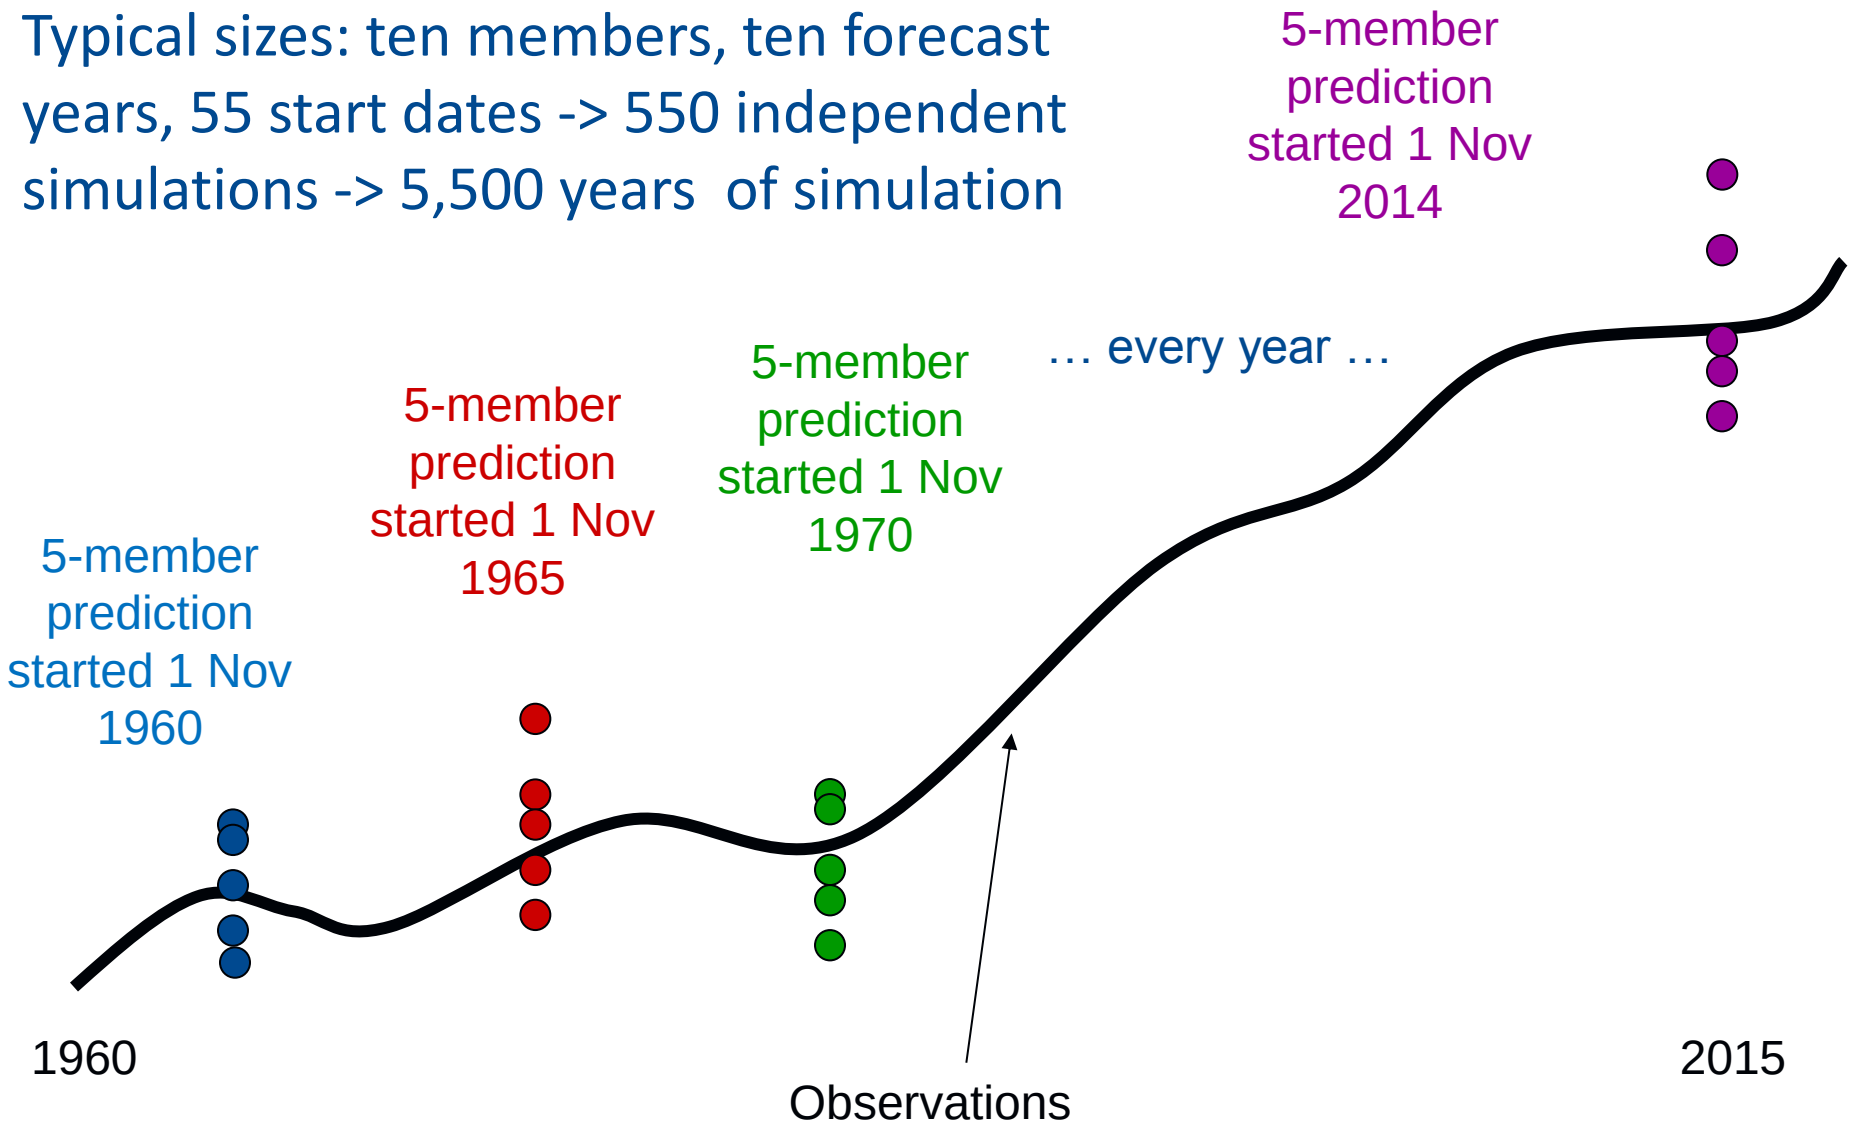

Predicting extremes

JJA near-surface temperature correlation of the ensemble mean from experiments with a climatological (top) and difference with one with realistic (bottom) land-surface initialisation. Results for EC-Earth2.3 started in May over 1979-2010.

Two ways for the analysis: online or offline

Objective: reducing data traffic

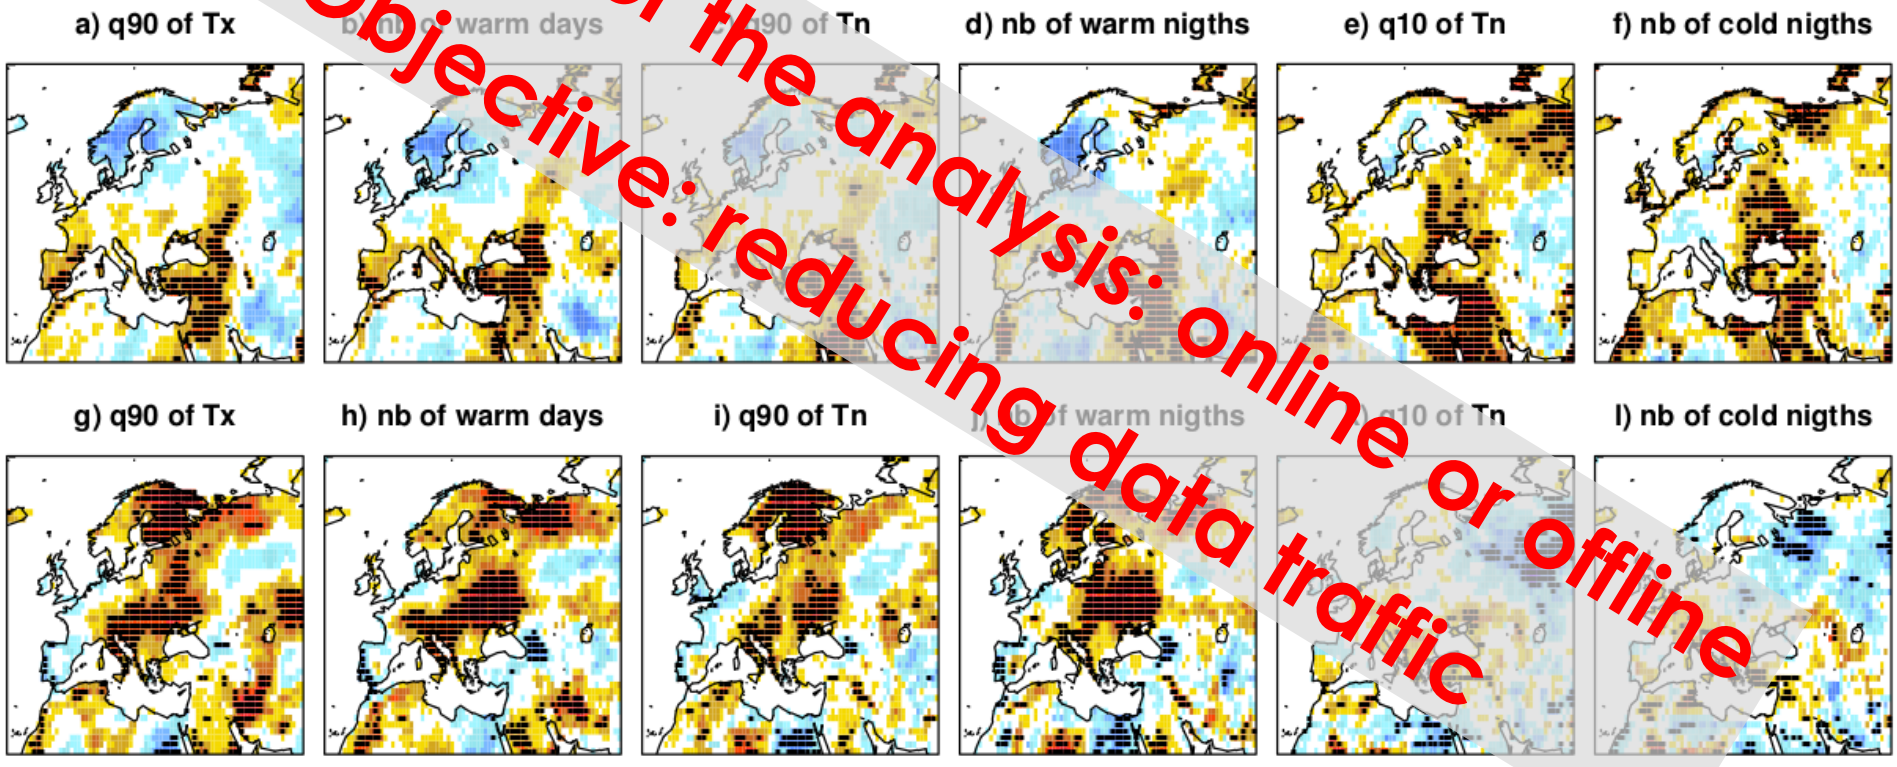

Prodhomme et al. (2015, Clim. Dyn.)

Attribution of Antarctic sea ice records

2014 was an exceptional year for the Antarctic sea-ice extent. A set of sensitivity experiments with NEMO allows to attribute it to anomalous southerly advection of cold air (Indian sector) and ocean pre-conditioning (Ross Sea).

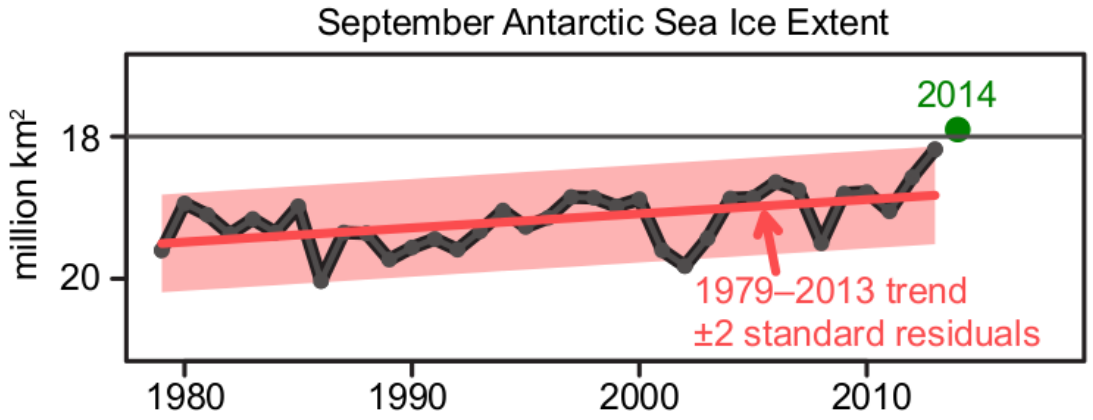

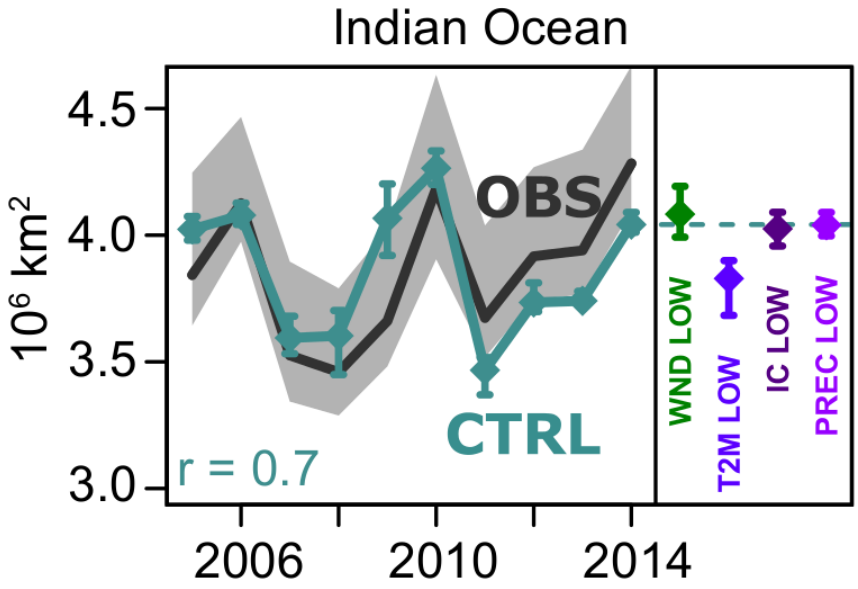

Massonnet et al. (2015, BAMS)

Decadal climate predictions

Global-mean near-surface air temperature and AMV against GHCN/ERSST3b for forecast years 2-5.

Global mean surface air temperature (GMST)

Atlantic multidecadal variability (AMV)

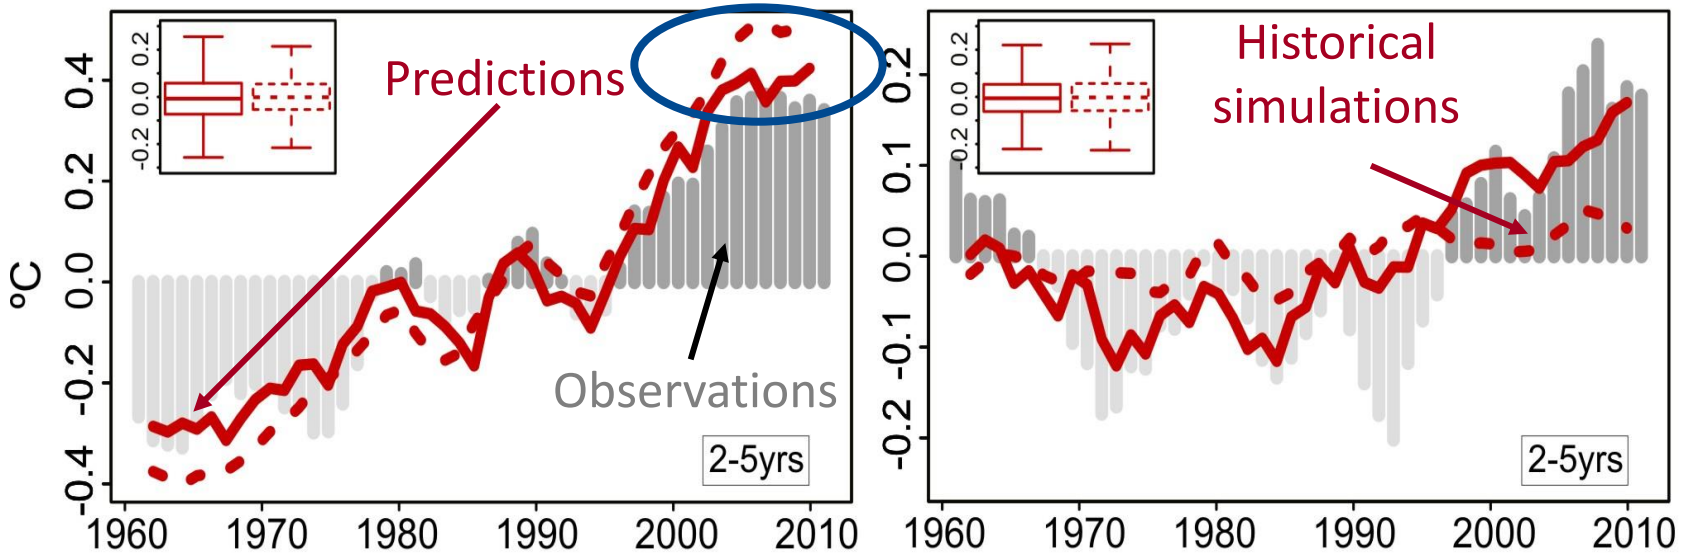

Initialised simulations reproduce the global temperature and some of the AMV tendencies and suggest that initialization corrects the forced model response and phases in internal variability.

Doblas-Reyes et al. (2013, Nat. Comms.)

Climate prediction is very expensive

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Climate prediction allows running jobs independently and simultaneously by wrapping together ensemble members for different start dates. This is not trivial parallelisation.

A suitable workflow manager is required.

5,500 simulated years
1.4 PB, 200 million core hours

EC-Earth3 at Lindgren, PDC						
Number of Start Dates		1	5	10	10	20
Number of Members		5	5	10	10	10
Number of Independent Simulations		1	25	50	100	200
T159-ORCA1	Cores	144	3600	7200	14400	28800
	Wall-clock Time (Hours) / year	5	5	5	5	5
	CPU Time (Hours) / year	720	18000	36000	72000	144000
	Output Size (GB) / year	10,80	480	960	1920	3840
T255-ORCA1	Cores	360	9000	18000	36000	72000
	Wall-clock Time (Hours) / year	5	5	5	5	5
	CPU Time (Hours) / year	1800	45000	90000	180000	360000
	Output Size (GB) / year	19,20	5184	10368	20736	41472
T799-ORCA025	Cores	1104	27600	55200	110400	220800
	Wall-clock Time (Hours) / year	40	40	40	40	40
	CPU Time (Hours) / year	44160	1104000	2208000	4416000	8832000
	Output Size (GB) / year	256,80	6420	12840	25680	51360

Computational Earth sciences

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Computational Earth Sciences

1. Ensure an efficient use of the computational resources by the research groups.
2. **Development of HPC user-friendly software framework for Earth system modeling and the management of operational systems.**
3. **Data management.**

Computer sciences beyond support

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Performance Team

- Provide HPC services such as performance analysis
- Apply new computational methods

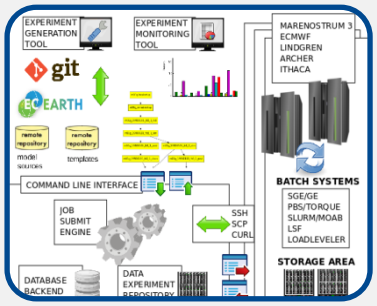

Models and Workflows Team

- Development of HPC user-friendly software framework
- Support the development of atmospheric research software

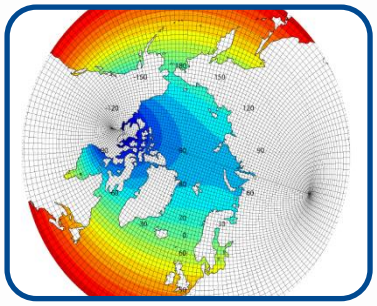

Data and Diagnostics Team

- Big Data in Earth Sciences
- Data services
- Visualization

Make the most of the expertise in data and computational technologies at BSC

A social challenge: energy use

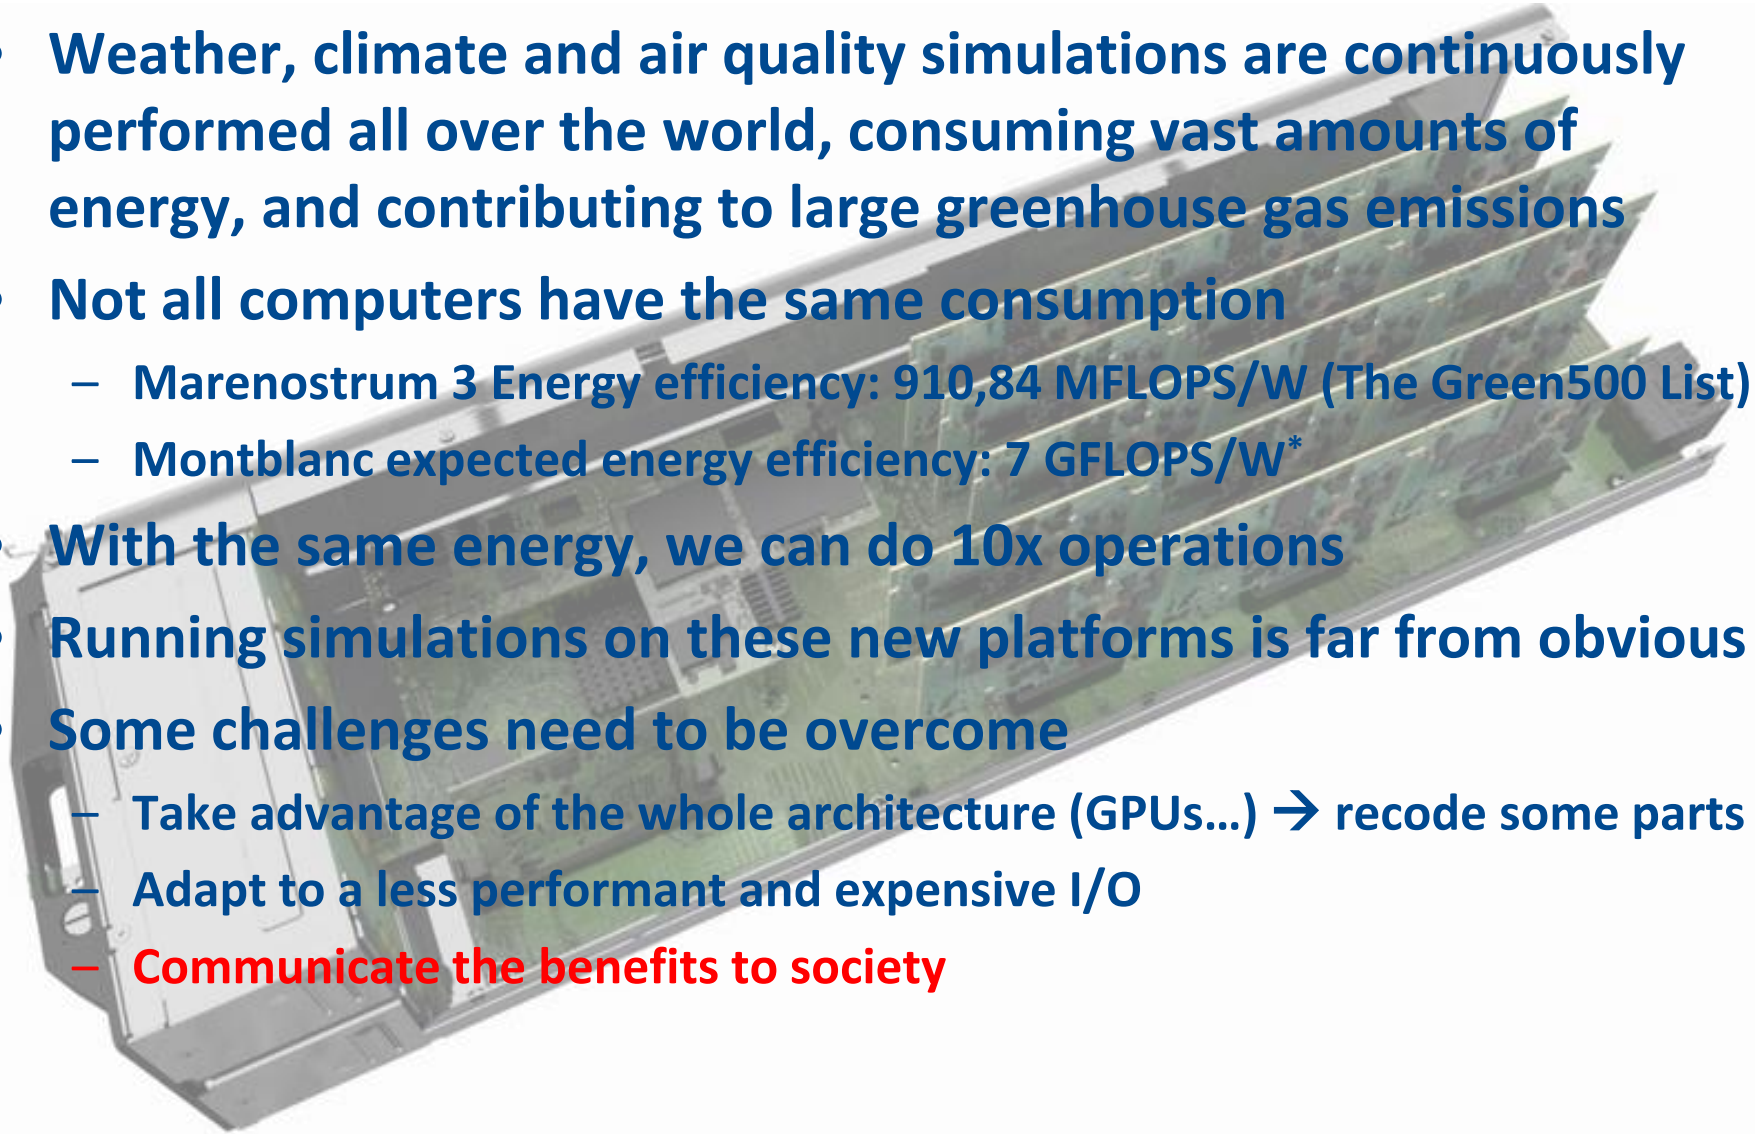

- **Weather, climate and air quality simulations are continuously performed all over the world, consuming vast amounts of energy, and contributing to large greenhouse gas emissions**
- **Not all computers have the same consumption**
  - **Marenostrum 3 Energy efficiency: 910,84 MFLOPS/W (The Green500 List)**
  - **Montblanc expected energy efficiency: 7 GFLOPS/W***
- **With the same energy, we can do 10x operations**
- **Running simulations on these new platforms is far from obvious**
- **Some challenges need to be overcome**
  - **Take advantage of the whole architecture (GPUs…) → recode some parts**
  - **Adapt to a less performant and expensive I/O**
  - **Communicate the benefits to society**

* http://www.exascale.org/mediawiki/images/6/6b/Talk11-Ramirez.pdf

Efficiency in Earth Science models

- Simulations use a **huge amount of resources**.
- Future simulations will need much more resources.
- NEMO: **Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**.
- Almost 170.000 lines of FORTRAN 90 code.
- Parallelization based on spatial domain decomposition through MPI.
- Mostly small stencil element calculations.
- A performance analysis shows that there are too many global communications and that only 20% of the time is spent in calculations.

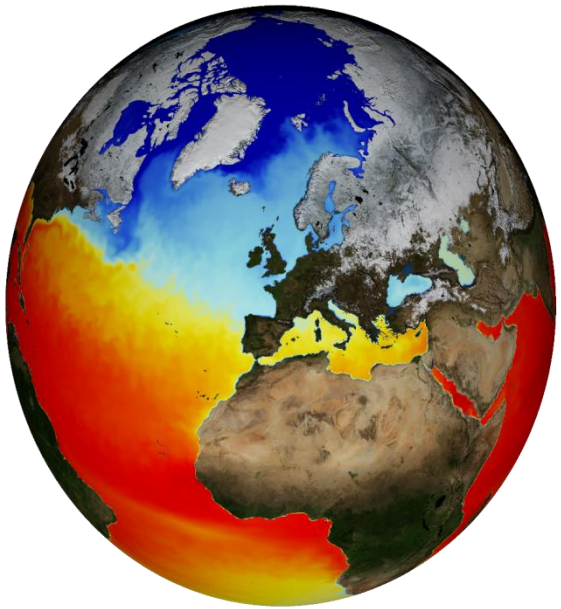

0,84 M points

ORCA 2
550 MB of memory
8 CPU hours
10 Gigabytes of output (daily)

68 M points

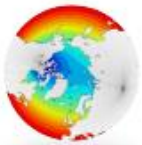

ORCA ¼
47 Gigabytes of memory
3500 CPU hours
120 Gigabytes of output (daily)

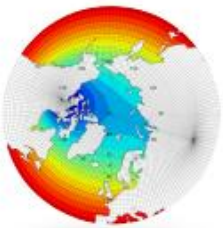

991 M points

ORCA 1/12
414 Gigabytes of memory
90 000 CPU hours
1 Terabyte of output (daily)

Ocean model optimisation

MPI message packing

Taking in account that NEMO is really sensitive to latency, messages aggregation is the best way to reduce the time invested in communications. Therefore, consecutive messages have been packed wherever the computational dependencies allow to do so.

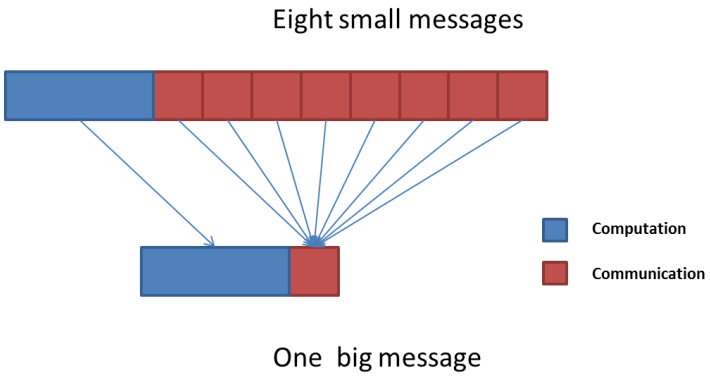

Convergence check reduction

Some routines use collective communications to perform a convergence check in iterative solvers. The cost of this verifications is really high, reaching a 66% of the time. Wherever the model allowed it, we reduced the frequency of this verifications in order to increase parallel efficiency.

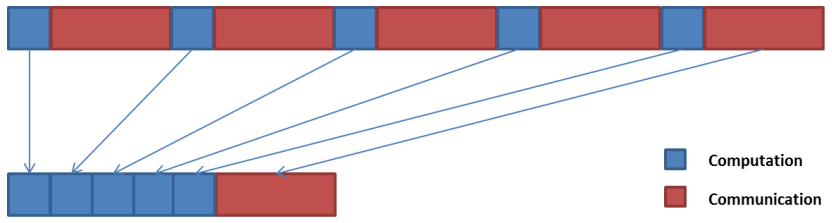

Reordering

In order to apply the message packing optimization to as many routines as it was possible, it was necessary to rearrange some computation and communication regions, taking into account the dependencies between them , to reduce the number of messages. This way it was possible to compute (and communicate) up to 41 variables at the same time, resulting in a dramatic reduction of the granularity.

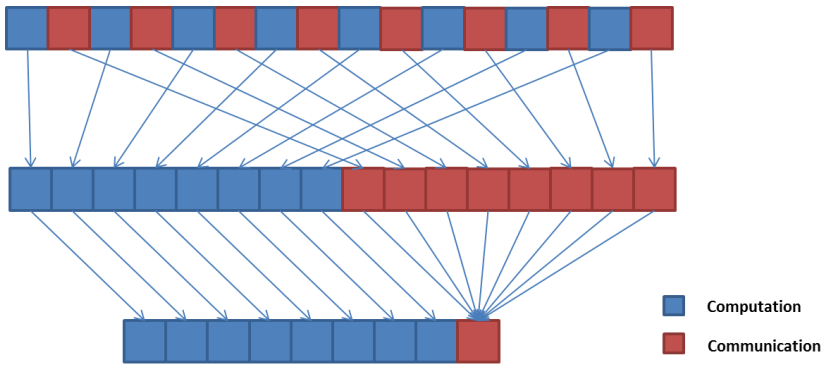

Ocean model optimisation

Original code

Optimized code

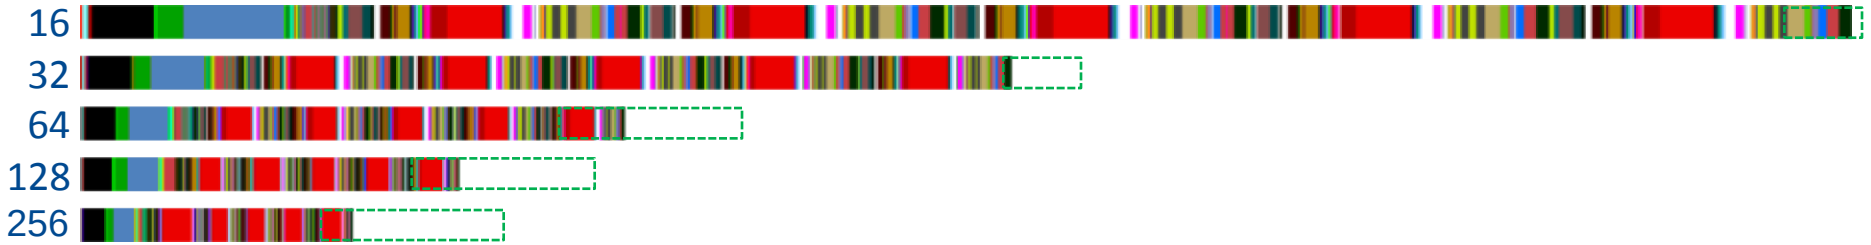

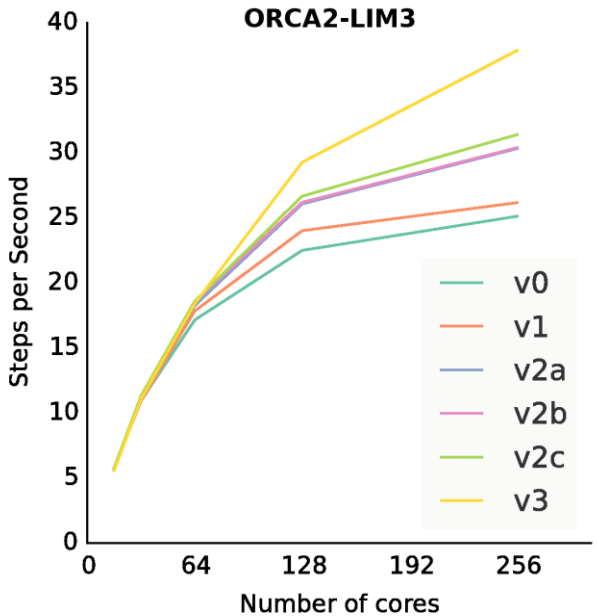

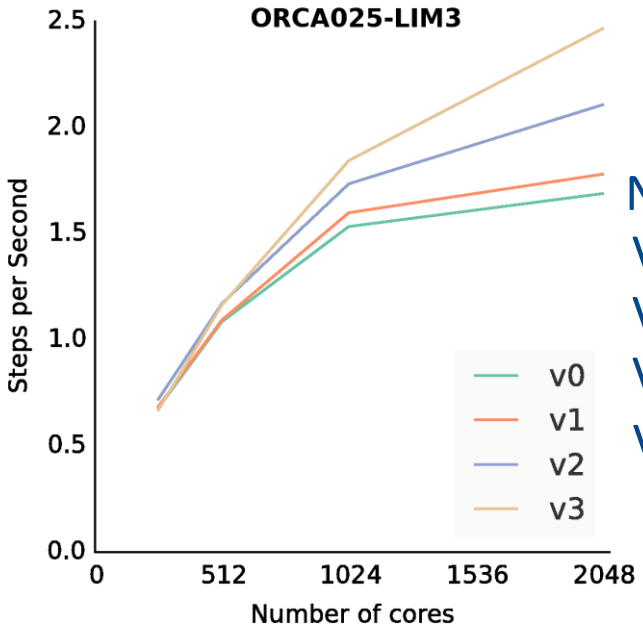

NEMO/LIM 3.6
V0 → Original
V1 → Message packing
V2 → Conv. check reduction
V3 → Reordering

EC-Earth optimisation

ESMs need coupling between components to communicate boundaries -> communication and interpolation. Coupling improvements:

- Message and calculations aggregation
- Parallelize transformation process

Optimizations improve the execution time for coupling from original version (v0) to the improved version (v1) up to 40%.

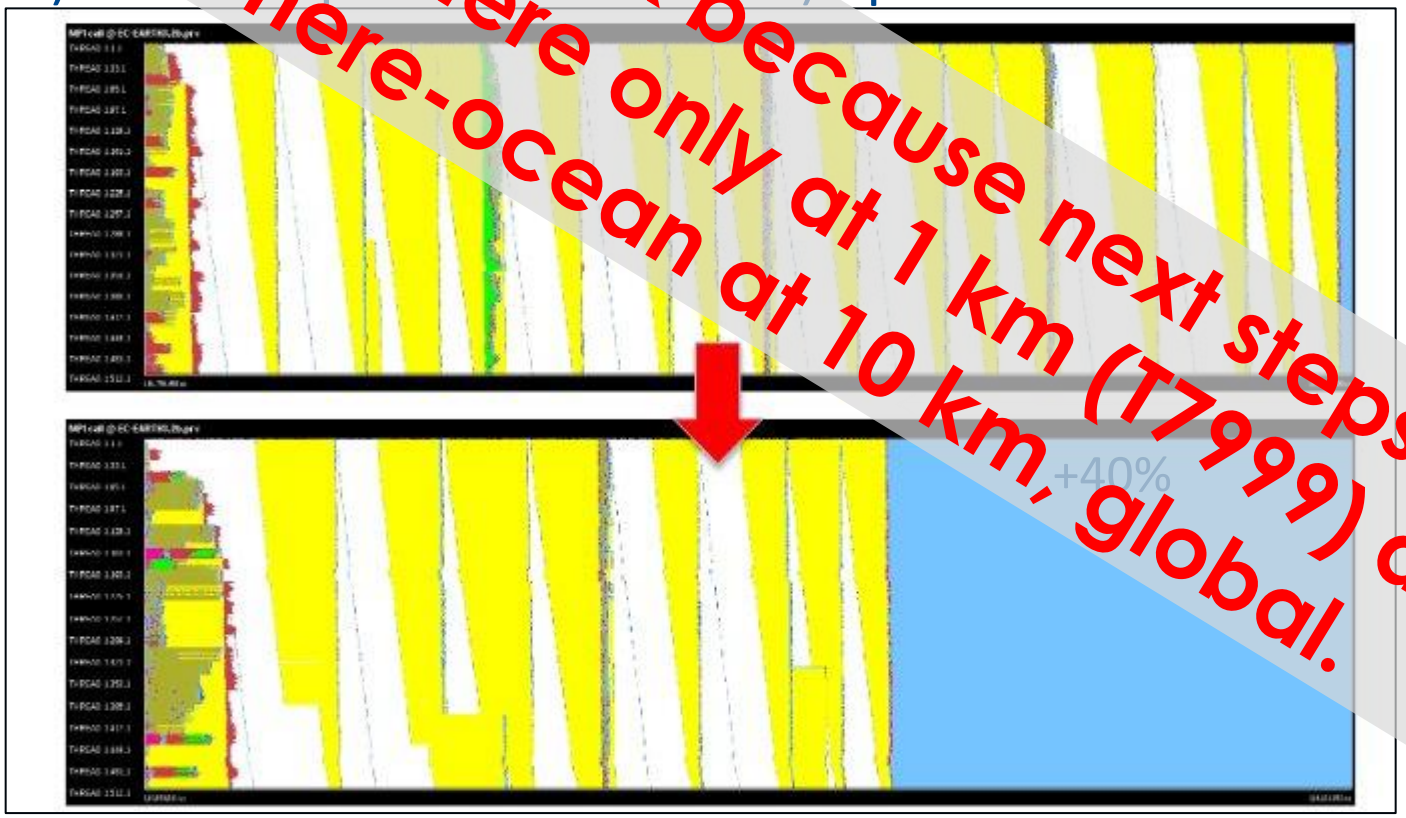

Fundamental work because next steps are to run atmosphere only at 1 km (T7999) and atmosphere-ocean at 10 km, global.

EC-Earth reproducibility

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Strike a balance between reproducibility, accuracy and performance.

Temperature difference of the ensemble means (five-members) between two experiments using FP_precise and O2. Black doted regions for statistically significant differences (K-S test).

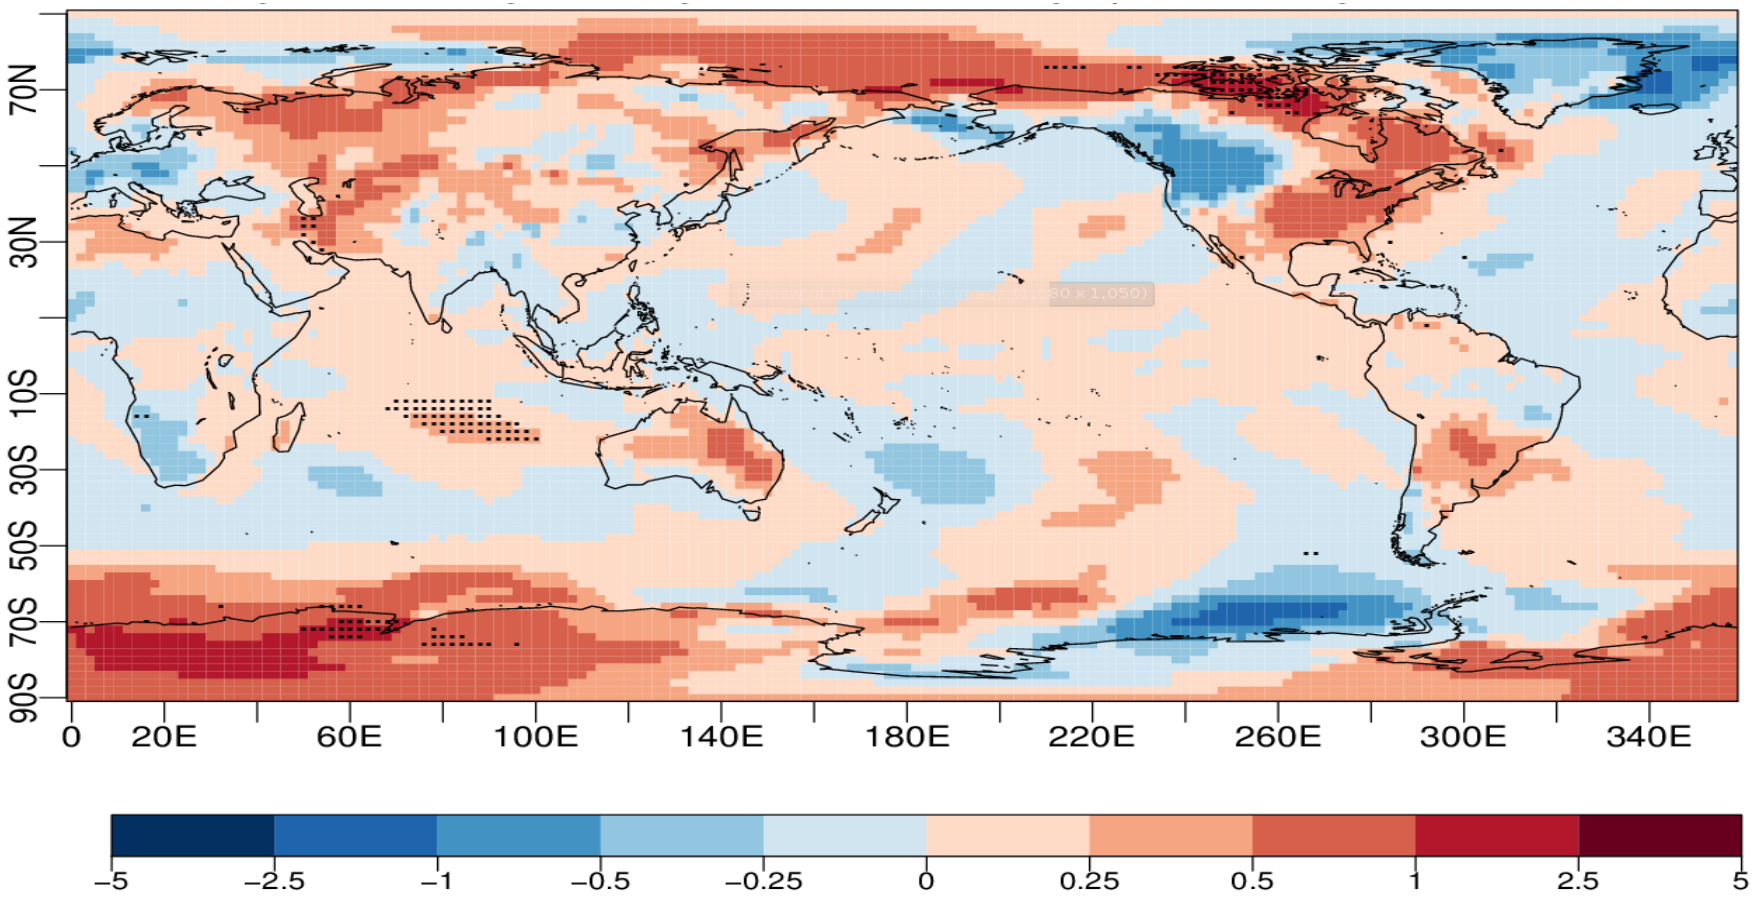

NMMB/BSC-CTM optimisation

NMMB: Non-hydrostatic Multiscale Meteorological Model on the B-Grid

- Meteorological core developed by NCEP
- Aerosols and gas-phase chemistry by BSC

Daily operational run in Marenostrum (16 nodes).

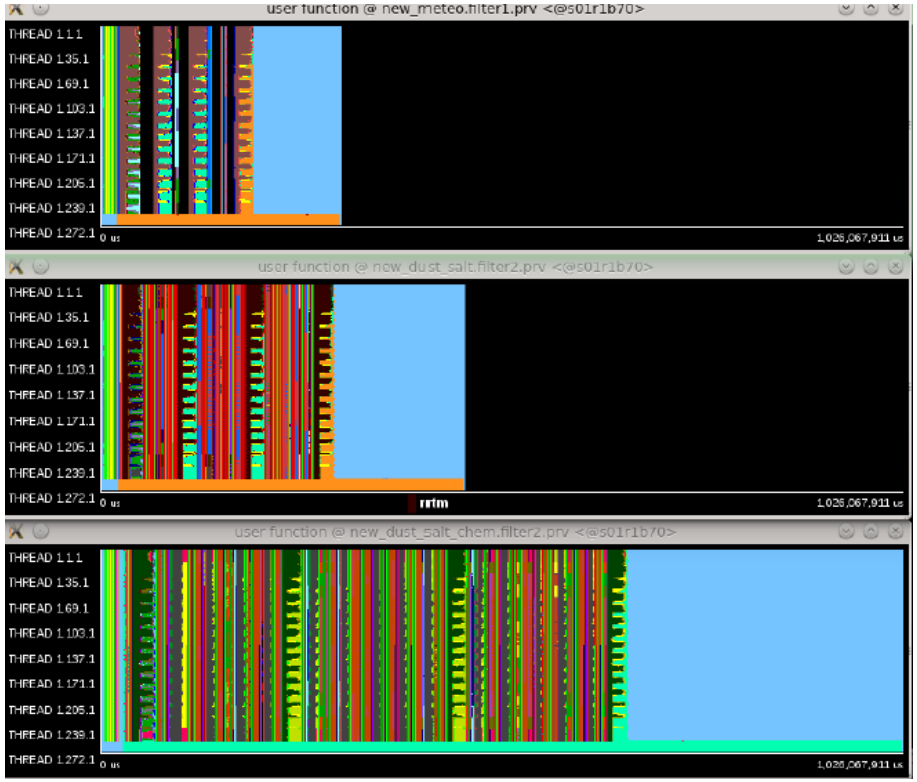

One hour simulation of NMMB/BSC-CTM, global, 24 km, 64 layers

meteo: 9 tracers

meteo+aerosols: 9+16 tracers

meteo+aerosols+gases: 9+16+53 tracers

NMMB/BSC-CTM workflow

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Data assimilation (EnKF), 12-member ensemble, use Autosubmit.

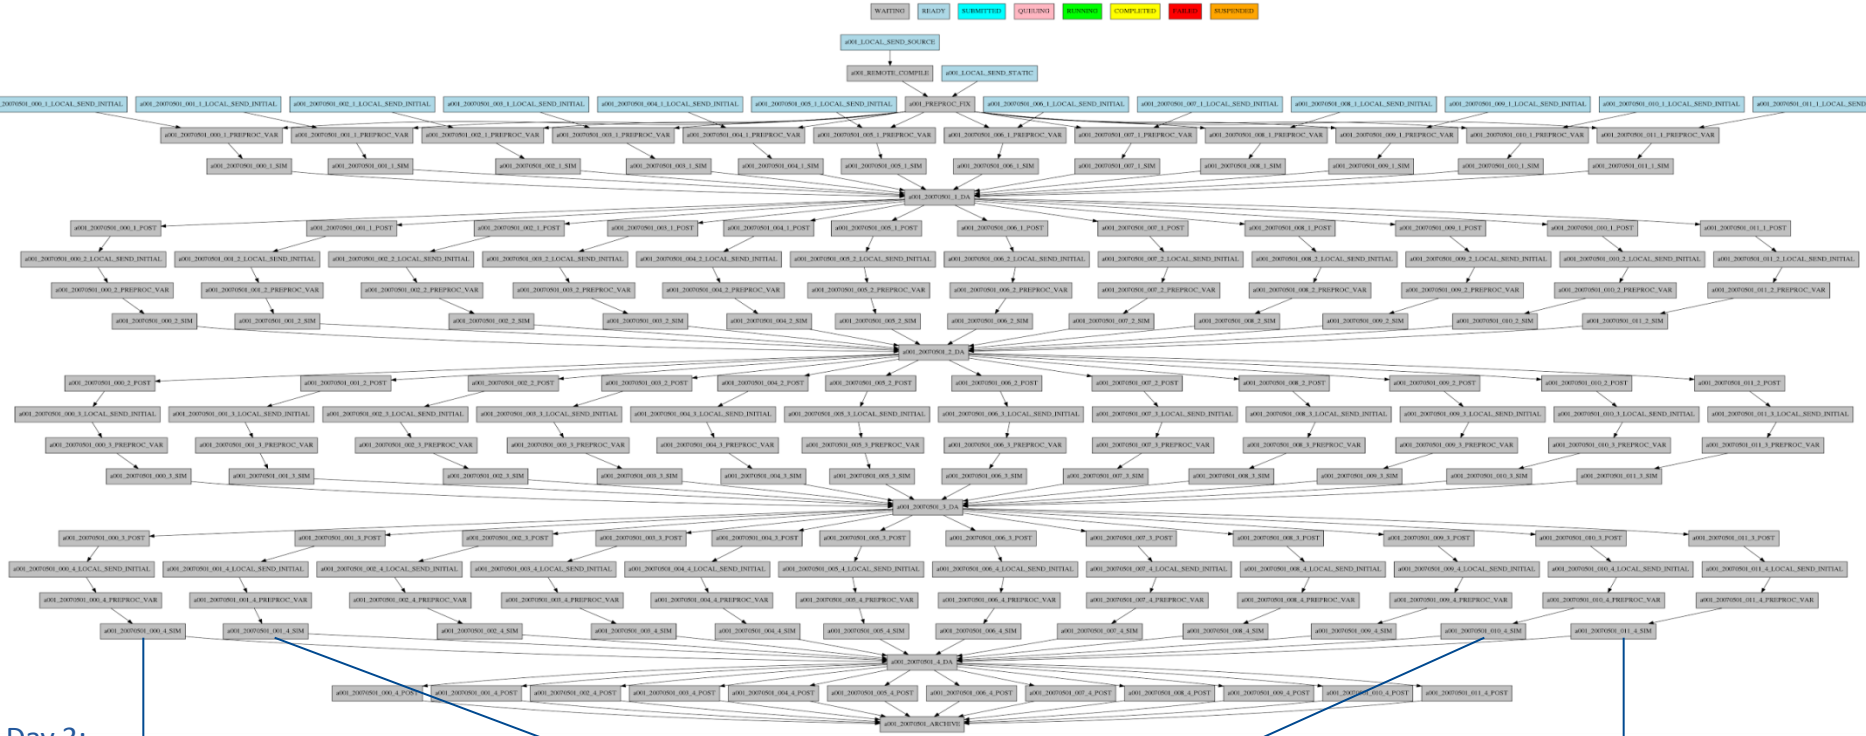

Day 3:

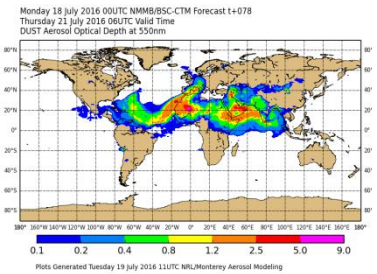

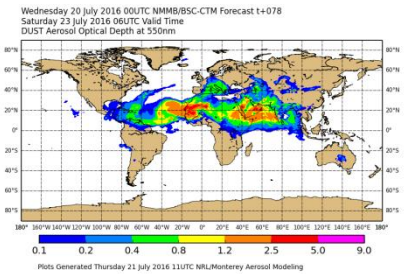

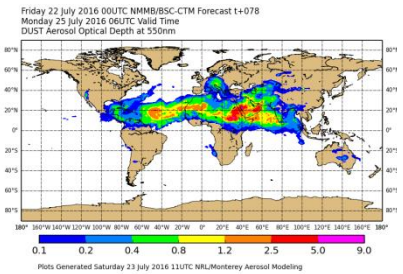

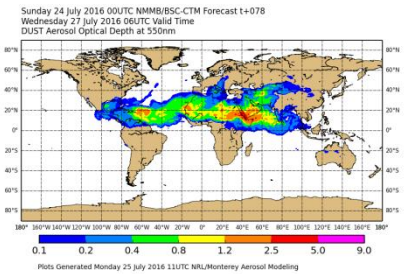

Online diagnostics

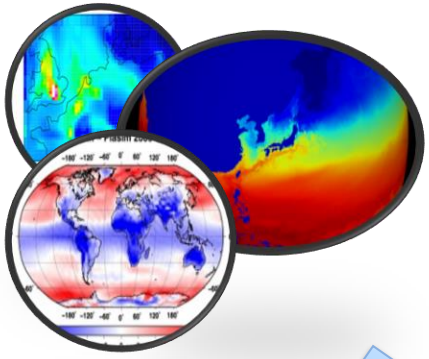

Climate model

Current approach: offline diagnostics

HPC → Model output → Serial node → Diagnostics ← End user

Our aim: online diagnostics

HPC → Model output → Diagnostics as part of the I/O ← End user

Analytics as a Service

- Diagnostics online during the run (either on compute node or elsewhere)
- Reduced data traffic and storage
- New diagnostics (data mining of extremes) possible
- The user gets the results faster → crucial to develop climate services (both public and private)

Earth system services

Earth System Services

1. Design of user-defined research to guide the fundamental science of the BSC-ES department.
2. **Development of user-interaction platforms to tailor and disseminate knowledge, tools and technology transfer.**
3. **Advancement of user and societal benefits via engagement, co-production and feedback from users to the provider.**

Service-driven research

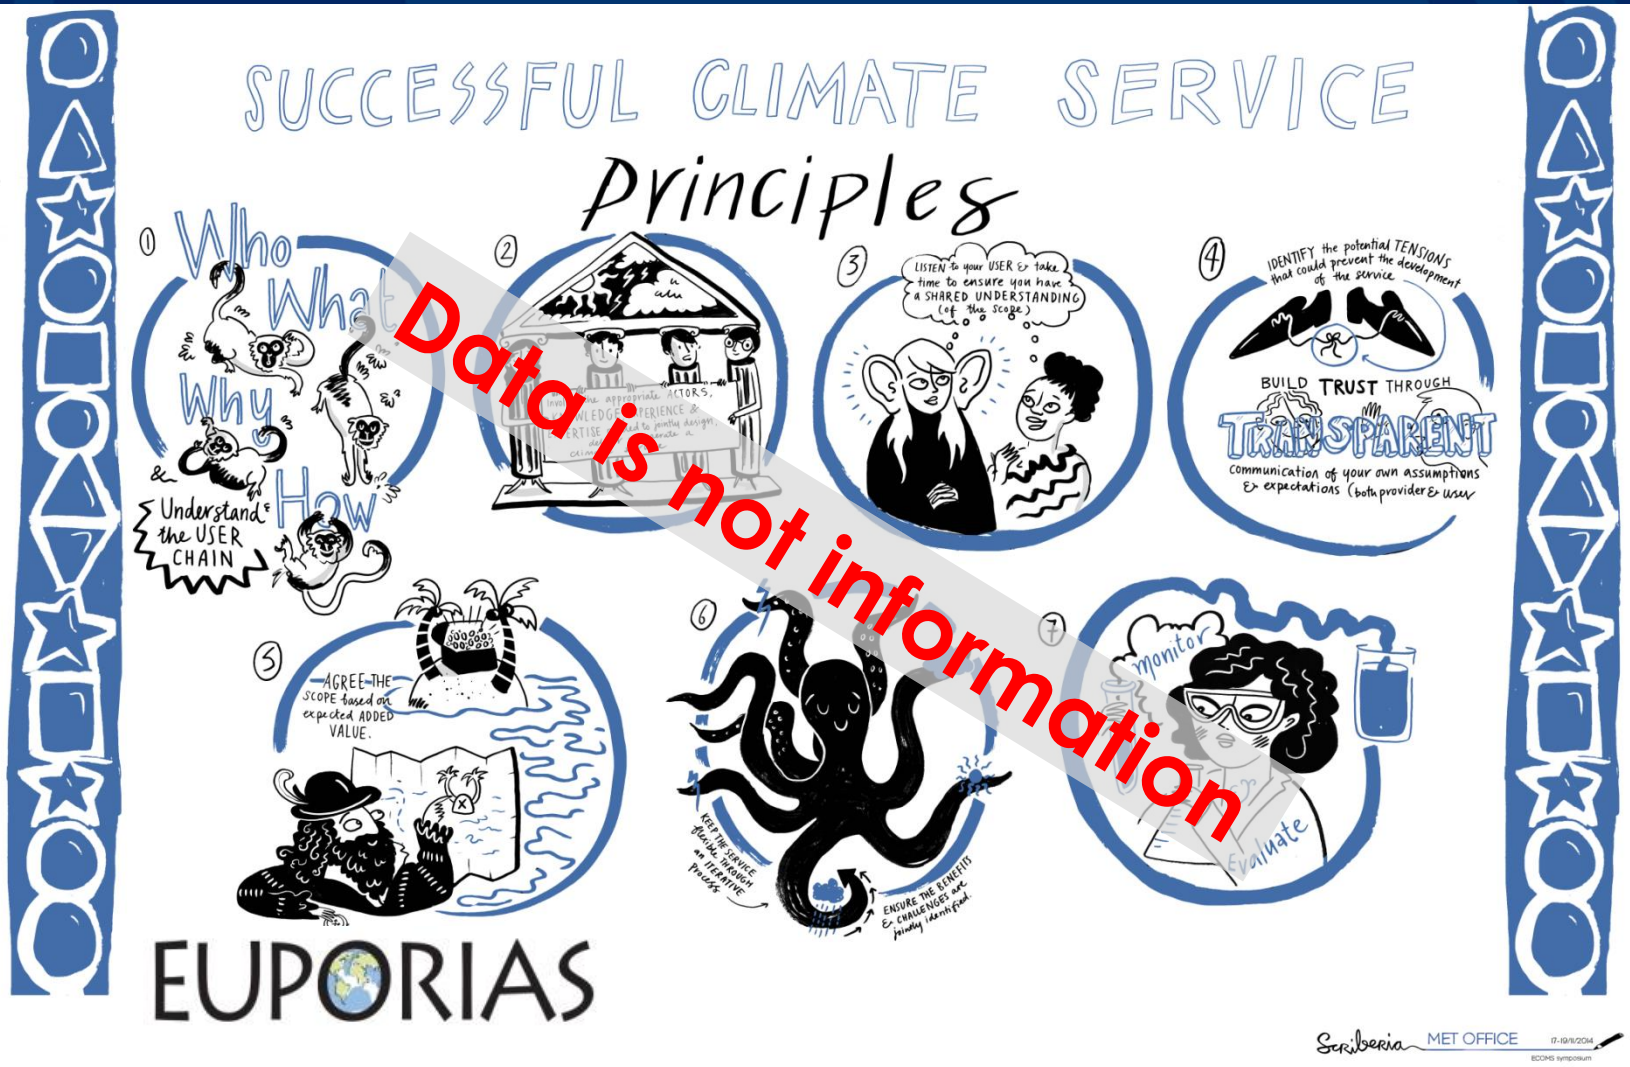

Ethical Framework for Climate Services four core elements: integrity, transparency, humility and collaboration.

Air quality forecasts for the public

Air quality forecast system: CALIOPE

Provides air quality related information for the coming days and for the application of short term action plans for air quality managers.

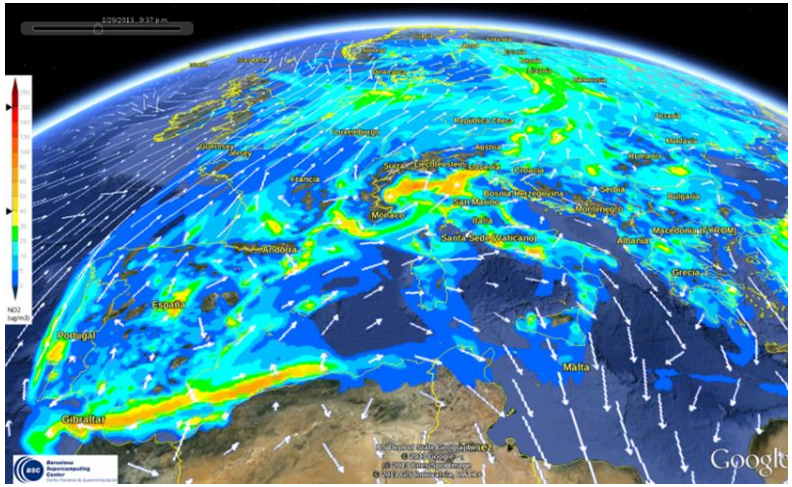

Information is delivered using both online or custom applications:

www.bsc.es/caliope

ANDROID APP ON Google play

Available on the App Store

Air quality impact assessment

Air quality modelling provides comprehensive description of air quality problems by relating emission sources and atmospheric conditions

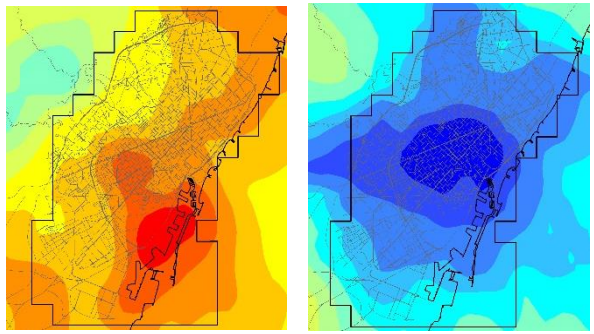

Left: NO_2 maximum h values in Barcelona (red >200 µg/m^3)
Right: Reductions due to fleet electrification (blue >25 µg/m^3)

Seasonal wind power predictions

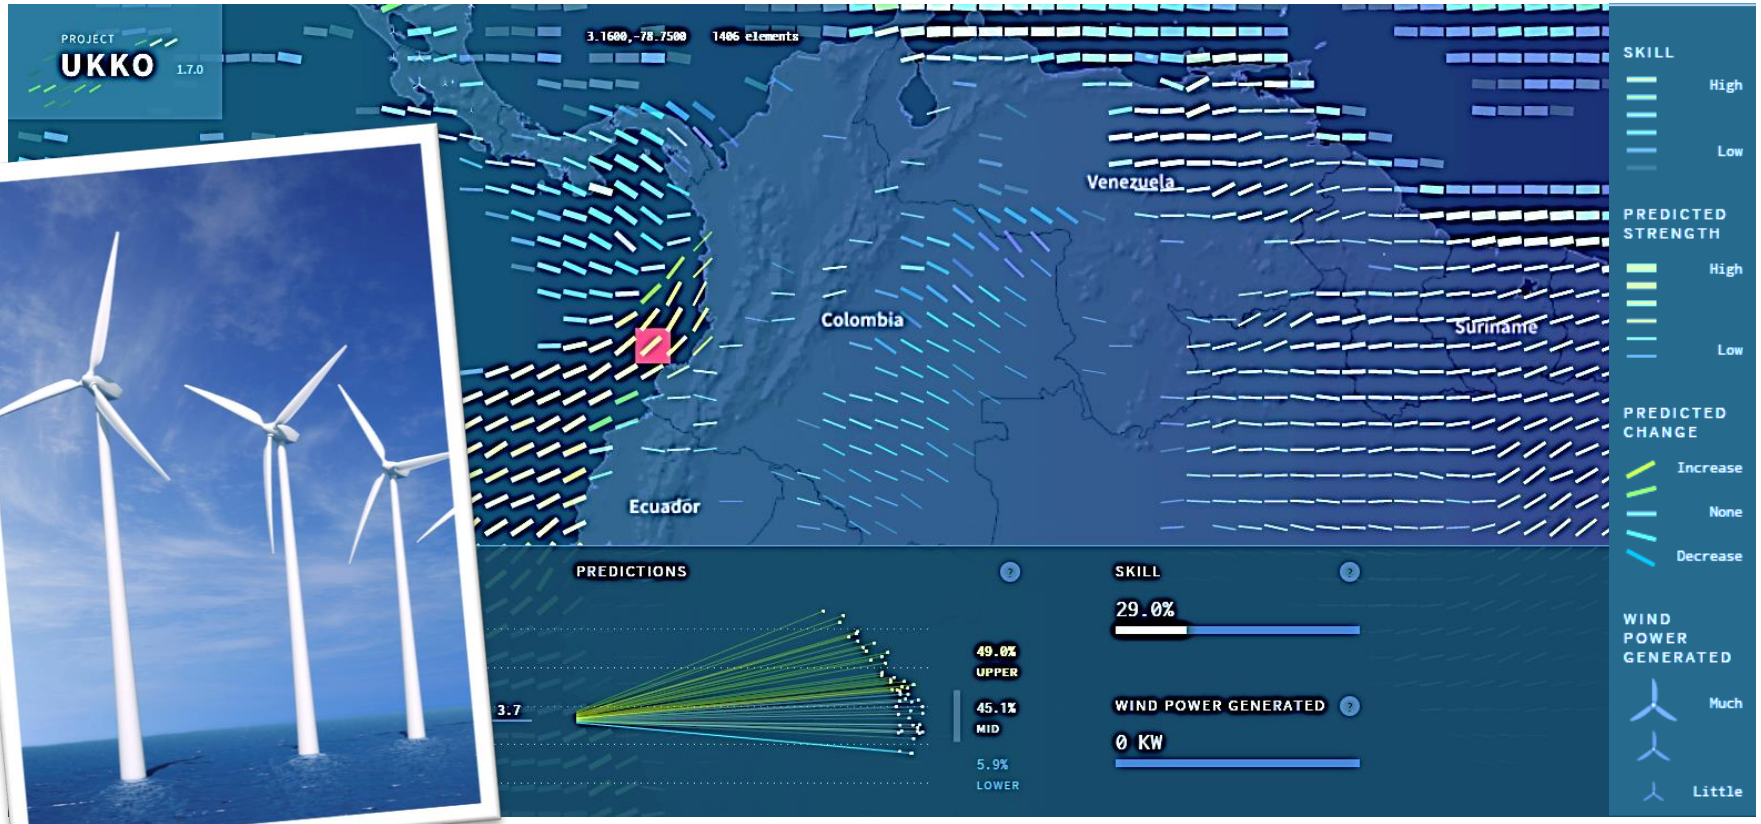

RESILIENCE
PROTOTYPE

SPECS
EUPORIAS
SEVENTH FRAMEWORK PROGRAMME
Clim4Energy
NEWA
EDF
edp renewables
ALSTOM
VORTEX
natural power
EnBW

Seasonal predictions of hurricanes

First comprehensive service of predictions of tropical cyclone seasonal frequency

http://www.bsc.es/ESS/seasonalhurricanepredictions/

Predictions/projections for agriculture

Testing seasonal forecast for MARS: BSC and JRC are exploring how the MARS Crop Yield Forecasting System (MCYFS) could ingest the seasonal forecast for a future operational use.

Current approach of MARS

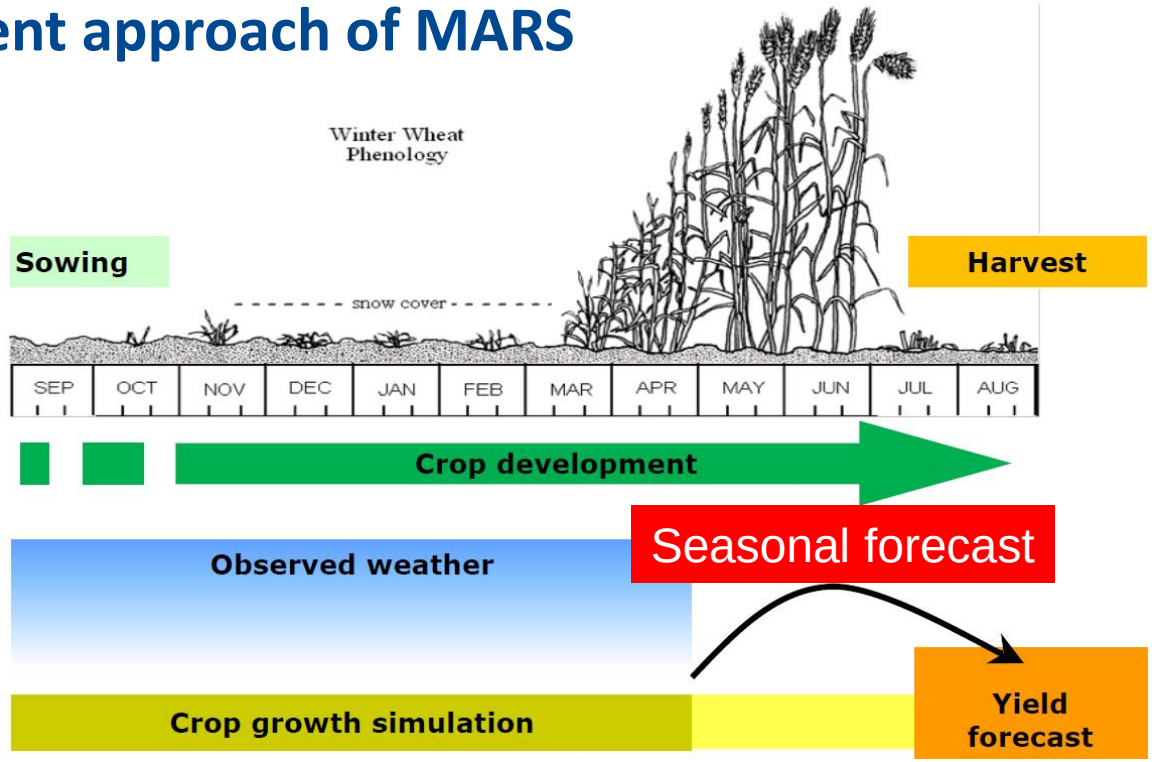

Some open questions

A new paradigm has come to stay: user-driven research

- **Education**: in the era of open data, take advantage of the open education opportunities.
- **Heterogeneity**: link to and merge our data with communities with larger impact (urban, arts, social).
- **Standards**: in a collaborative environment standards are a must and everyone's responsibility.
- **Technology**: make the most of a rapidly evolving technology (heterogeneous nodes, data software, mobile data capture, storage/compression, computing and storage outsourcing, visualisation).
- **Industry engagement:** how can we solve the problem of involving more efficiently the private sector?